arXiv:1410.4382v2 [q-fin.RM] 19 Nov 2015

VERIFICATION OF INTERNAL RISK MEASURE ESTIMATES

MARK H.A. DAVIS

Abstract. This paper concerns sequential computation of risk measures for financial data and asks how, given a risk measurement procedure, we can tell whether the answers it produces are 'correct'. We draw the distinction between 'external' and 'internal' risk measures and concentrate on the latter, where we observe data in real time, make predictions and observe outcomes. It is argued that evaluation of such procedures is best addressed from the point of view of probability forecasting or Dawid's theory of 'prequential statistics' [Dawid, JRSS(A)1984]. We introduce a concept of 'calibration' of a risk measure in a dynamic setting, following the precepts of Dawid's weak and strong prequential principles, and examine its application to quantile forecasting (VaR – value at risk) and to mean estimation (applicable to CVaR – expected shortfall). The relationship between these ideas and 'elicitability' [Gneiting, JASA 2011] is examined. We show in particular that VaR has special properties not shared by any other risk measure. Turning to CVaR we argue that its main deficiency is the unquantifiable tail dependence of estimators. In a final section we show that a simple data-driven feedback algorithm can produce VaR estimates on financial data that easily pass both the consistency test and a further newly-introduced statistical test for independence of a binary sequence.

JEL classification: C13 C32 C44 C53 G17
Key words: Risk measures, probability forecasting, prequential statistics, quantile and mean forecasting, consistency of estimates.

1. Introduction

Computing risk measures is a matter of primary importance to the financial services industry, both from the point of view of short-term risk management and for regulatory capital allocation purposes; see Embrechts and Hofert (2014) for a recent survey. For a portfolio of assets, a risk measure is generally interpreted as some functional of the conditional distribution F of the portfolio loss[1] between times t and $t+h$ given all the market information up to today, time t. The prediction horizon h is typically a week or 10 days for market risk management, somewhat longer for credit-related assets or insurance. The most widely used risk measures are the *value at risk* VaR, and CVaR, variously known as *conditional value at risk*, *expected shortfall* or *expected loss beyond* VaR. Formal definitions are as follows. The definition of CVaR for general distributions involves some subtleties, for which the reader is referred to Rockafellar and Uryasev (2002).

Definition 1.1. *Let F be a right-continuous distribution function on the real line and $\beta \in (0,1)$.*

Date: November 20, 2015.
The author thanks Paul Embrechts for introducing him to this topic, Phil Dawid for discussions of prequential statistics, and Johanna Ziegel for extremely helpful comments and advice. He is also grateful to two anonymous referees for their careful reading of the paper and many constructive suggestions for improvements. Some of this work was carried out at the Hausdorff Research Institute for Mathematics at the University of Bonn, where the author was a visitor during the Trimester Program *Stochastic Dynamics in Economics and Finance* in August 2013. Discussions there have been very helpful in clarifying the content of this paper.

[1]In the literature on risk management there is no settled convention whether losses should be designated as positive or negative. In this paper losses are positive, so we are concerned with the *right-hand* tail of the distribution. Distribution functions will be taken to be right-continuous, i.e. $F_Y(y) = \mathbb{P}[Y \le y]$.

(i)The β-quantile *is the interval* $[q_\beta^-, q_\beta^+)$ *if* $F(q_\beta^+) > F(q_\beta^+-)$ *and the interval* $[q_\beta^-, q_\beta^+]$ *otherwise, where* $q_\beta^- = \inf\{y : F(y) \geq \beta\}$ *and* $q_\beta^+ = \inf\{y : F(y) > \beta\}$.
(ii) The value at risk at level β *is* $\mathrm{VaR}_\beta = q_\beta^-$.
(iii) CVaR_β *is the mean of the* β*-tail distribution, given by*

$$F_\beta(y) = \begin{cases} 0, & y < q_\beta^- \\ \frac{F(y)-\beta}{1-\beta}, & y \geq q_\beta^-. \end{cases}$$

Explicitly,

$$\mathrm{CVaR}_\beta = \mathrm{VaR}_\beta + \frac{1}{1-\beta}\int_{(q_\beta^-,\infty)} (y - q_\beta^-)F(dy) \tag{1.1}$$

$$= \frac{1}{1-\beta}\left[\int_{F(q_\beta^-)}^1 q_\tau^- d\tau + q_\beta^-(F(\beta)-\beta)\right]. \tag{1.2}$$

Expression (1.1) quantifies the gap between VaR and CVaR, while (1.2) provides a relationship between CVaR and quantiles. When there is no jump at the β-quantile, (1.2) reduces to the familiar expression

$$\mathrm{CVaR}_\beta = \frac{1}{1-\beta}\int_\beta^1 \mathrm{VaR}_\tau d\tau.$$

There has been renewed debate about the relative merits of VaR and CVaR and indeed about the risk management process as a whole. A huge literature on risk measures was triggered off by the seminal paper by Artzner, Delbaen, Eber, and Heath (1999) in which axioms for a 'coherent' risk measure were formulated, but much of this literature is presented in a pure mathematical framework taking no account of where the data is coming from, how the risk measure is to be computed or what the ultimate purpose of the exercise is. Clearly these practical considerations have to be included in any evaluation of the risk management process. For example, Cont, Deguest, and Scandolo (2010) concentrate on stability of computation under perturbations of the model and conclude that computation of CVaR suffers from unavoidable instabilities that are not present in the computation of VaR, challenging the at the time conventional wisdom that CVaR is to be preferred because it is coherent while VaR is not. In a similar vein, Kou, Peng, and Heyde (2013) draw the distinction between 'external' and 'internal' risk management, which we discuss in Section 2.3 below. They give a revised set of axioms appropriate to external risk management and show that value at risk does satisfy these axioms. They make the interesting suggestion of replacing CVaR by CMVaR, the conditional *median* shortfall. Of course, at threshold level β, CMVaR is just VaR at level $\frac{1}{2}(1+\beta)$, so CMVaR satisfies the Kou-Peng-Heyde axioms and has all the computational advantages of VaR while making some attempt at quantifying tail risk. Further comments will be found in Section 7.

In this paper we take a different tack. Given that we have selected a risk measure and a computational algorithm, how can we tell whether the answer is 'correct' when we apply the algorithm to real data? This is by no means a simple question, for one very clear reason. In evaluating, say, the value at risk at level β, what we are computing at time $(k-1)$ is $(F_k^m)^{-1}(\beta)$, the βth quantile of the conditional distribution F_k^m of the portfolio return at time k given all information up to $(k-1)$, computed according to a chosen model, labelled m. Even if the model m is time-invariant, F_k^m is a *different distribution* for each k, because the conditioning event is different. When time k arrives, we observe *one number*, which may or may not exceed the predicted quantile level. How to evaluate the quality of such predictions is the province of *probability forecasting* (Dawid, 1986, Lai, Shen, and Gross, 2011, Gneiting, 2011), a branch of statistics that until very recently was largely ignored by researchers in financial risk

management, although some of its techniques are routinely used in various ways by practitioners. The approach we take is inspired by P. Dawid's theory of *prequential probability* (Dawid, 1984), and in particular to the exceptionally stimulating paper Dawid and Vovk (1999).

The paper is laid out as follows. In Section 2 we discuss the fundamentals of time series prediction in general terms, highlighting the special features of financial price data. We also introduce the essential distinction between *external* and *internal* risk measures together with some comments on the applicability of the prequential approach in the latter case. Section 3 defines elicitability from statistical decision theory and summarises the information we need. The core of the paper is Section 4, in which Definition 4.2 formalizes the concept of *calibration* of a risk management statistic in a dynamic setting, and we demonstrate its connection to elicitability via identification functions. In Section 5 we examine the case of quantile forecasting and VaR estimation in some detail and show that this case has especially favourable features: the quantile statistic is calibratable under essentially no conditions on the underlying data (Theorem 5.2). In this section we also introduce an auxiliary test for independence, the details of which are given in Appendix A, and we also discuss an example, given by Holzmann and Eulert (2014), showing that calibration and independence, while necessary conditions for correct prediction, are far from sufficient. Section 6 covers mean estimation, including CVaR. The main calibration result here is Theorem 6.5, utilising martingale convergence theorems. We discuss the CVaR problem in more detail in Section 7; first we describe the Rockafellar and Uryasev (2000) characterisation of CVaR as the solution to a minimisation problem and its role as a further verification test for VaR predictors; then we highlight the main problem in this area: unquantifiable tail dependence. Finally, we illustrate these ideas by looking at one week ahead quantile prediction for the FTSE100 stock index, showing that a simple data-driven algorithm can produce a sequence of quantile forecasts that easily survive calibration and independence tests.

2. Fundamentals of prediction

2.1. **Background.** Financial risk management is essentially about prediction: given whatever information we think relevant, we have to take a view on the likely returns of a portfolio over some holding period, generally quantified by some risk measure. It may be helpful to start out by placing this problem in the general context of time series prediction. In any such problem, the approach taken must depend on the nature of the data, on what it is we are trying to predict and on what the purpose of the prediction is. There is a hierarchy of possibilities.

(i) The simplest case is coin tossing: the whole probabilistic structure is fixed axiomatically, with no need for statistical modelling. No-one will quarrel with the statement that the distribution of the number of heads in the next n tosses is the Binomial distribution $B(n, 0.5)$.

(ii) There are situations in which the data is produced by a well-understood and stable physical mechanism, for example Geiger counter detection of radioactive emissions. Here it is clear from the physics that the sequence of counts will constitute a Poisson process. We can estimate its rate from past data, and there is no reason to suppose that the rate will be different in the future, or at most will vary in predictable ways, at least over short time scales.

(iii) Next, weather forecasting. This is similar to the previous case (ii) in that predictions are mainly derived from mathematical models describing the underlying physics, but of course the latter are extremely complex; see Warner (2010) for an authoritative account. Prediction is bound up to a great extent with studies of the effects on the models of perturbations in model parameters and/or initial conditions, a subject that has become a discipline in its own right under the name of *Uncertainty Quantification* (Smith, 2014). Other approaches include the

more statistical topic of pattern matching, in which periods in the past are identified in which the weather patterns match those of today.

It is useful to distinguish two very different problems in climate science (a) short-term weather forecasts (up to a week) and (b) extreme-value problems such as flood barrier design, estimates of the probability of inundation of coastal areas or incidence and severity of tornados. Short-term forecasts can be monitored by checking whether the predictions are well 'calibrated'. We give an example below, and indeed the main purpose of this paper is to formalize this idea. Techniques to deal with problems in case (b) are completely different because, by definition, the data is sparse. Flood barrier design is an engineering problem in which complex mathematical models and extreme-value theory (Embrechts, Klüppelberg, and Mikosch, 1997) are combined to give best-possible estimates of the frequency of exceedance of various high-water marks; this is followed by a cost-benefit analysis in which the level of protection is traded off against the cost of providing it. The point is that we are never in the position of case (a) where a sequence of predictions can be checked against subsequent outcomes. Instead, physics and data are combined to assemble evidence in favour of a decision the consequences of which will, we hope, never be tested to destruction.

(iv) Statistics comes into its own in situations where we have no physical model but an adequate supply of data that is 'reasonably predictable'. A case in point would be a sequence of insurance claims resulting from car accidents. There is no physical theory, but there is a huge amount of relevant data. Individual accidents are largely independent, and the general prevalence of accidents and the claims arising from them depend on well-understood factors such as the age distribution of drivers, repair costs, growth in traffic and improvements in safety due to engineering developments, speed restrictions, etc. In these circumstances highly credible statistical models can be built, giving an accurate view of the claims likely to be faced by an insurance company.

2.2. **Financial risk management.** Now we move to the subject of this paper, the computation of risk management parameters for portfolios of financial assets. The problem could hardly be more different from those in cases (i)-(iii) above. The underlying reason for this is that the economy is a meso-scale phenomenon: too big to be modelled in complete detail, but too small, and too interconnected, to be treated by methods of statistical mechanics.

As a representative data set we will take the series displayed in Figure 2.1(a), 30 years of weekly values S_n of the FTSE100 stock index 1984-2013. The accompanying Figure 2.1(b)

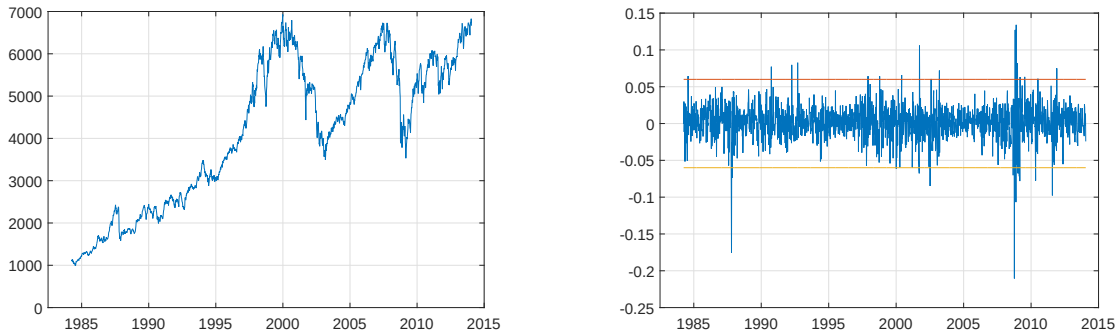

(a) Index values. (b) Returns.

FIGURE 2.1. FTSE100 index: weekly values 1994-2013

shows the associated series of returns $Y_n = (S_n - S_{n-1})/S_{n-1}$ and demonstrates the typical stylised features found in financial price data: apparent non-stationarity and highly 'bursty'

volatility. The horizontal rules are at levels ±0.06, approximately the 1% and 99% quantiles. These will be needed in Section 7.1.

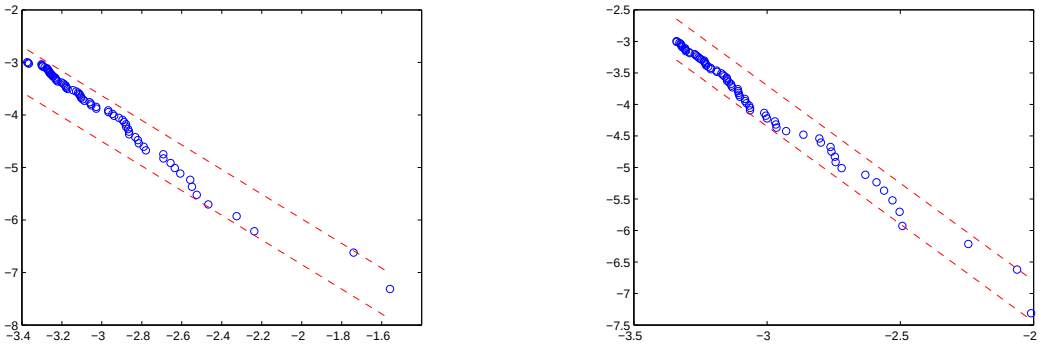

(a) Left tail (x-axis reversed) (b) Right tail

FIGURE 2.2. Empirical distribution, left and right 5% on log-log scale.

Figure 2.2 shows the left and right 5% =50 points of the empirical return distribution on a log-log scale. Based on this data, we conclude that the empirical distribution has power law tails with indices $\kappa = 2.35, 3.25$ respectively [2]. Of course, these series have been the subject of intensive research over at least the last 50 years. A notable—and perhaps the most original—contributor was Benoît Mandelbrot (Mandelbrot and Taylor, 1967, Mandelbrot, 1997), who introduced the heavy-tailed *fractional Brownian motion* as an asset price model, and the subject has become mainstream in econometrics (see for example Bollerslev, 1986, Campbell, Lo, and MacKinley, 1990) and statistics (Cox, Hinkley, and Barndorff-Nielsen, 1996). An excellent account from a 'quant' perspective is Cont (2001). It is however a curious fact that remarkably little of this effort has been aimed directly at prediction[3], although from the perspective of contemporary risk management little else matters.

2.3. **External risk measures.** We mentioned above the distinction drawn by Kou et al. (2013) between 'external' and 'internal' risk measures. External risk measures are those employed by industry regulators in procedures imposed across the board on all regulated institutions, while internal risk measures are those used in individual institutions, or even trading desks, for day-to-day monitoring of the risks of trading books.

The external/internal distinction mirrors exactly the distinction between cases (b) and (a) respectively of weather forecasting, discussed in Section 2(iii) above. External risk measures are part of the process, depicted in Figure 2.3, by which the regulator imposes capital charges on the bank in order to provide an adequate cushion against trading losses (Basel Committee on Banking Supervision, 2013). Models must be built to compute return distributions F_k, $k = 1, \dots, n$, under various scenarios, and then the capital charge C is equal to some function of $\{\mathfrak{s}(F_1), \dots, \mathfrak{s}(F_n)\}$, where $\mathfrak{s}$ is the VaR or CVaR at some very high level such as 99.5% or 99.75%. The input data D may be taken from the immediate past or/and from 'stressed' periods in history, and is used for model calibration. Having calibrated the model, computations are invariably done by Monte Carlo simulation, so the computed F_k *always* have finite support; this point is stressed by Kou et al. (2013) and by Rockafellar and Uryasev (2002) and is relevant to our discussion of elicitability below.

[2]A distribution function F on $\mathbb{R}$ has power left tail with index $\kappa > 0$ if $F(x) \sim |x|^{-\kappa}$ as $x \to -\infty$. To estimate κ from an empirical tail sample we find, for a given κ, the tightest $0 < c_1 < c_2$ such that $c_1|x|^{-\kappa} \le F(x) \le c_2|x|^{-\kappa}$ for all x in the sample and then minimise c_2/c_1 over κ, giving the bounds shown in Figure 2.2. An analogous procedure applies to the right tail.

[3]There is an elegant paper on prediction of fractional Brownian motion by Gripenberg and Norros (1996).

There is no conceptual issue relating to the meaning of the return distributions F_k since they are outputs from a well-defined stochastic model. The whole process is an engineering job exactly analogous to flood barrier design. The important question in this area is numerical stability: we want to avoid a situation where widely differing capital charges are imposed on different banks merely because of minor variations in their internal models; see Cont, Deguest, and Scandolo (2010) for an excellent study of this problem. A much bigger question, of course, is whether the decomposition implied by Figure 2.3 of the map $D \mapsto C$ is the best way to arrive at an appropriate capital cushion. There are dissenting voices, see for example Haldane (2012).

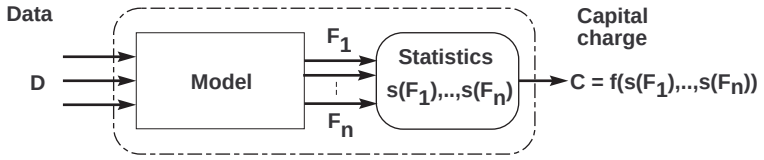

FIGURE 2.3. Capital charge allocation process.

In this paper we focus on *internal* risk management—predicting the risks faced by trading books. Here the quantile level is typically much lower, say 95%, so we can expect to see occasional exceedances and can monitor their frequency. This is the kind of problem probability forecasting is designed to address.

2.4. **Falsifiability.** First, let us consider a foundational issue. We compute what we claim to be the conditional distribution F of future returns and/or some statistic $\mathfrak{s}(F)$. But in this universe of highly non-stationary data, and given that no resampling is possible, we might well ask whether the predictive distributions implied by statistical models have any meaning at all. A useful reference point is the falsifiability test of Karl Popper (2002): a statement is *meaningful* if and only if it is *falsifiable*, i.e. evidence could in principle be produced that would show the statement to be false[4]. Now consider the following statement $\mathfrak{S}$: 'the conditional distribution of the FTSE100 return Y_k, given data up to time $k-1$, is F', where F is a specified distribution function. According to Popper's criterion, statement $\mathfrak{S}$ is surely meaningless. We compute F at time $k-1$, and at time n we get a single number $Y_k = x$; so was F correct? $\mathfrak{S}$ is falsified at time k only if x lies outside the support of F, which will never be the case in practice, where the support is invariably specified as $\mathbb{R}$ (or $\mathbb{R}^+$ for long-only portfolios). Since subsequent data points $Y_{k+1}, Y_{k+2}, \ldots$ are drawn from different conditional distributions, they cannot be said to provide much useful evidence about the correctness of F, and in any case *post hoc* data is not germane, since decisions have to be made on the basis of calculations at time $k-1$ and history cannot be rewritten afterwards. Consequently $\mathfrak{S}$ is not falsifiable, meaning that any statement about F must depend on uncheckable *a priori* modelling assumptions.

What is needed here is a shift of perspective. Instead of asking whether our model is correct, we should ask whether our objective in building the model has been achieved. This view of the prediction problem is standard in in some other areas of science, and was in fact pioneered in connection with weather forecasting (see for example Joliffe and Stephenson, 2003). An example will illustrate the point, taken from Dawid (1986). In many countries it is customary for weather forecasters to predict the probability of rain the next day in quantised form $0.0, 0.1, 0.2, \ldots, 1.0$. The obvious way to evaluate such forecasts is, for each $n = 0, \ldots, 10$, to calculate over time

[4]The concept is related to the basic asymmetry between proof and counterexample: to show $A \Rightarrow B$ we have to show that in *every* case where A holds, B holds too, whereas to show $A \not\Rightarrow B$ we only have to find *one* case where A holds but B does not.

the relative frequency of rain on days after the forecast probability was $n/10$. We then plot the relative frequencies against the predicted probabilities to obtain a *reliability diagram*. Nothing is or can be asserted about the accuracy of the forecast on *any particular* day.

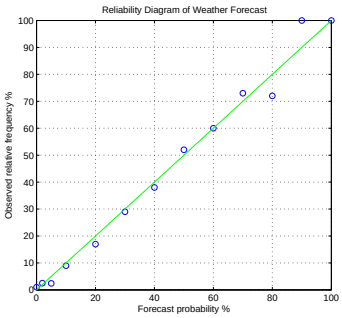

FIGURE 2.4. Reliability diagram for Chicago forecaster.

Figure 2.4 shows the results of 2820 12-hour forecasts by a forecaster in Chicago in the period 1972-76 (Dawid, 1986); no-one can be in any doubt that this forecaster was doing a good job: the forecasts are well-calibrated. A key point here is that the evaluation accords with the principles of P.Dawid's 'prequential' theory of statistics (Dawid, 1984). These principles, as enunciated in Dawid and Vovk (1999) are

Weak prequential principle: Evaluation of forecasting systems should be based only on the observed data and the numerical values of the forecasts produced (not on the algorithm that produced them).

Strong prequential principle: Criteria for correct prediction should only depend on agreement between Nature and Forecaster on the stochastic law $\mathbb{P}$ generating the data, not on what that law is (within some specified class $\mathcal{P}$).

Although formal application of these principles barely figures in the published literature on risk management, related methods are universally applied in the industry under the name of 'back-testing'[5] A large part of the 'VaR vs. CVaR' debate is concerned with the question whether it is true that if a statistic is not elicitable then it cannot be back-tested, see Acerbi and Szekely (2014) for a recent contribution. While we do not settle this question here, we do provide some formal structure within which the question can formulated in more precise terms.

3. Elicitability

The initial motivation for writing this paper was the striking set of results obtained by Gneiting (2011) and Ziegel (2014) on the elicitability properties of VaR and CVaR. These authors showed that CVaR is not elicitable, and this was used in various quarters as an argument against its use as a risk-management statistic in place of VaR, which is elicitable. This argument has now fallen by the wayside, as Fissler and Ziegel (2015) have recently shown that the pair (VaR, CVaR) is *jointly* elicitable, but nonetheless the controversy brought something new and important to the world of risk management. The circle of ideas relates to a decision-theoretic framework whose origins go back at least to work by L.J. Savage (1971), but the elicitability concept itself is due to Osband and Reichelstein (1985) and the name was coined

[5]This term is perhaps misleading as it seems to imply some special programme to re-live history rather than a procedure that is part of day-to-day practice. 'Monitoring' would be a better description.

by Lambert, Pennock, and Shoham (2008). The reader can consult Gneiting (2011) for a wide-ranging exposition of this subject.

We consider the probability space $(\mathbb{R}, \mathcal{B}, \mathbb{P})$, where $\mathcal{B}$ is the Borel σ-field. Y will denote the identity function $Y(y) = y \in \mathbb{R}$ and as usual the probability measure $\mathbb{P}$ is identified with the (right-continuous) distribution function of Y. It is a familiar fact that if $Y \in L_2(\mathbb{R}, \mathcal{B}, \mathbb{P})$ then the function $f(x) = \mathbb{E}[(x - Y)^2]$ achieves its minimum at $x = \mathbb{E}[Y]$ and this is true whatever the distribution F within the L_2 class. Elicitability is concerned with generalizing this characterization of the mean value to other statistics $\mathfrak{s}(F)$ of the distribution function. For a given statistic $\mathfrak{s}(F)$, can we find a *score function* $S(x, y)$ such that $x \mapsto \mathbb{E}_F[S(x, Y)] = \int S(x, y)F(dy)$ is minimized at $x = \mathfrak{s}(F)$ for all F in some wide class $\mathcal{F}$ of distributions? In general $\mathfrak{s}(F)$ may be set-valued, as is the case for the *β-quantile*, and *a fortiori* for the median, equal to $q_{\frac{1}{2}}$.

Our choice of score function will, as in Gneiting (2011), be restricted to measurable functions $S : \mathbb{R}^2 \to \mathbb{R}$ satisfying

(i) $S(x, y) \geq 0$ with equality if $x = y$

(ii) For each $y \in \mathbb{R}$ the function $x \mapsto S(x, y)$ is continuous, and is continuously differentiable if $x \neq y$.

We say that S is a *consistent scoring function for a statistic* $\mathfrak{s}$ relative to a class $\mathcal{F}$ of distribution functions F if whenever $Y \sim F \in \mathcal{F}$

$$\mathbb{E}_F[S(t, Y)] \leq \mathbb{E}_F[S(x, Y)] \qquad \forall\ t \in \mathfrak{s}(F), x \in \mathbb{R}. \tag{3.1}$$

S is *strictly consistent* if it is consistent and equality in (3.1) implies $x \in \mathfrak{s}(F)$.

Definition 3.1. *A statistic $\mathfrak{s}$ is* elicitable for $\mathcal{F}$ *if there exists a strictly consistent scoring function S.*

The attractive feature of the approach is the precision of the results: it is possible to prove mathematically, in relevant cases, that particular statistics *are* or *are not* elicitable. What is not so clear is how to apply these results in a dynamic context such as risk management where the data is a sequence $Y_1, Y_2, \ldots$ of random variables each having a different conditional distribution, particularly in view of the fact that the criterion (3.1) fails to respect the weak prequential principle. The next section suggests an answer to this question, but first we consider a few examples, following Gneiting (2011, §3).

3.1. **Examples.**

3.1.1. *Mean value.* Here $\mathcal{F}$ is the set of distributions with finite variance and the score function $S(x, y) = (x - y)^2$ is continuously differentiable. We can characterize optimality by noting that

$$\frac{\partial}{\partial x}\mathbb{E}[S(x, Y)] = \mathbb{E}\left[\frac{\partial}{\partial x}S(x, Y)\right] = x - \mathbb{E}[Y], \tag{3.2}$$

confirming that the expected score is indeed minimized at the mean value $\mathbb{E}[Y]$. $S = (x - y)^2$ is not the only score function eliciting the mean value—others exist that do not require the existence of second moments; see Section 3.1.3 below and Gneiting (2011, §3) for details and further examples.

3.1.2. *Quantiles and VaR.* Here $\mathcal{F}$ is the set of all probability distributions on some interval $I \subset \mathbb{R}$. Then the β-quantile, $\beta \in (0, 1)$ is elicitable. If I is compact then a score function S satisfying conditions (i), (ii) above is strictly consistent for the β-quantile if and only if it takes the form

$$S(x, y) = (\mathbf{1}_{(x \geq y)} - \beta)(g(x) - g(y)) \tag{3.3}$$

where g is a strictly increasing function. Score functions S as in (3.3) are strictly consistent without the compactness assumption in the class of distributions for which the random variable $g(Y)$ is integrable. An obvious choice is $g(y) = y$, but if we take g to be bounded and strictly increasing then no integrability condition is required.

Suppose g is continuously differentiable and let $\mathcal{F}_c$ be the class of continuous distribution functions. Then S is continuously differentiable except at $x = y$ and

$$\frac{\partial S}{\partial x} = g'(x)[\mathbf{1}_{(x\geq y)} - \beta]. \tag{3.4}$$

Since the event $(Y = x)$ has probability 0 for all $F \in \mathcal{F}_c$ we see that

$$\mathbb{E}\left[\frac{\partial}{\partial x}S(x, Y)\right] = g'(x)[F(x) - \beta], \tag{3.5}$$

which is equal to zero if and only if x is in the β-quantile set. If we drop continuity of the distribution function then

$$\frac{d}{dx}\mathbb{E}[S(x, Y)] = g'(x)[(1 - \beta)F(x-) + \beta F(x^+) - \beta]. \tag{3.6}$$

The expression on the right is negative if $x < q_\beta^-$ and positive if $x > q_\beta^+$, confirming that $\mathbb{E}[S(x, Y)]$ is minimized at any x in the β-quantile.

The value at risk VaR_β picks out one element, q_β^-, from the quantile set. Consequently, VaR is elicitable only in the set $\mathcal{F}_\uparrow \subset \mathcal{F}$ of *strictly increasing* distribution functions, for which the quantile set is a singleton.

3.1.3. *Expectiles.* For $\tau \in (0, 1)$ and $F \in L_1$ the τ-expectile is the unique solution m_τ to the equation

$$\tau \int_{(x,\infty)} (y - x)F(dy) = (1 - \tau) \int_{(-\infty,x)} (x - y)F(dy).$$

If ϕ is a C^1 strictly convex function, the score function

$$S(x, y) = (\tau \mathbf{1}_{(x<y)} + (1 - \tau)\mathbf{1}_{(x\geq y)})(\phi(y) - \phi(x) - \phi'(x)(y - x))$$

is strictly consistent for the τ-expectile in the class of F such that Y and $\phi(Y)$ are F-integrable. The natural choice is $\phi(x) = x^2$ when $(\phi(y) - \phi(x) - \phi'(x)(y - x)) = (y - x)^2$. If $\phi \in C^2$ then

$$\frac{\partial S}{\partial x} = \phi''(x)[\tau \mathbf{1}_{(x<y)} + (1 - \tau)\mathbf{1}_{(x\geq y)}](x - y), \tag{3.7}$$

and hence

$$\mathbb{E}\left[\frac{\partial}{\partial x}S(x, Y)\right] = -\phi''(x)\left[\tau \int_{(x,\infty)} (y - x)F(dy) - (1 - \tau)\int_{(-\infty,x)} (x - y)F(dy)\right]$$

so that $\mathbb{E}[(\partial S/\partial x)(S(x, Y))] = 0 \Leftrightarrow x = m_\tau$. This characterization only requires Y to be F-integrable. Note that the mean is the $\frac{1}{2}$-expectile, so a by-product here is a range of possible alternative score functions for the mean.

3.2. **Identifiability.** Given a class of distributions $\mathcal{F}$, an *identification function* for a statistic $\mathfrak{s}$ is a measurable function $V : \mathbb{R}^2 \to \mathbb{R}$ such that the expectation $\mathbb{E}[V(x, Y)]$ is well-defined whenever $Y \sim F \in \mathcal{F}$ and $x \in \mathbb{R}$, and

$$\mathbb{E}_F[V(x, Y)] = 0 \quad \Leftrightarrow \quad x \in \mathfrak{s}(F).$$

A statistic $\mathfrak{s}$ is *$\mathcal{F}$-identifiable* if an identification function exists. It is clear from (3.5) above that if $\mathfrak{s}$ is elicitable with score function S then under sufficient regularity conditions $V(x, y) = \partial S/\partial x(x, y)$ is an identification function. There is a kind of converse to this result, known as 'Osband's principle' (Gneiting, 2011, §2.4) according to which score functions can be obtained

from identification functions by a natural integration procedure. The relationship between the two is examined in detail by Steinwart, Pasin, Williamson, and Zhang (2014); their Corollary 9 asserts that for a scalar, single-valued statistic, under certain conditions elicitability is equivalent to existence of a bounded identification function, but the conditions include existence of a dominating measure, which is generally too restrictive.

3.3. **Distributional forecasts.** One way to predict a statistic is to predict the whole distribution and then calculate the statistic of the predicted distribution. This is the situation when one builds a stochastic model for the process of interest: past data is used to estimate model parameters and then it is a purely computational problem, tackled by analytic methods or by simulation, to evaluate the predicted distribution at some time in the future. Distributional forecasts may be evaluated by the use of *proper scoring rules* (Gneiting and Raftery, 2007, §2). If $\mathfrak{P}$ denotes the set of probability measures, or equivalently the set of distribution functions, on $(\mathbb{R}, \mathcal{B})$ then we can define $(\mathfrak{P}, \mathcal{B}_{\mathfrak{P}})$ as the Borel space corresponding to the topology of weak convergence. A *forecast* is a choice of $F \in \mathcal{F} \subset \mathcal{B}_{\mathfrak{P}}$ where $\mathcal{F}$ is some designated set of distributions. A scoring rule is a measurable function $\tilde{S} : \mathcal{F} \times \mathbb{R} \to \mathbb{R}$ such that the function $y \mapsto \tilde{S}(F, y)$ is G-integrable for all $G \in \mathcal{F}$. We define $\mathbf{S}(F, G) = \int \tilde{S}(F, y)G(dy)$. Here F is the prediction and G the 'true' distribution. The scoring rule[6] is *proper* if $\mathbf{S}(G, G) \leq \mathbf{S}(F, G)$ for all $F \in \mathcal{F}$ and *strictly proper* if $\mathbf{S}(F, G) = \mathbf{S}(G, G)$ implies $F = G$. If $S(x, y)$ is a score function for elicitation of a statistic $\mathfrak{s}$ then $\tilde{S}(F, y) = S(\mathfrak{s}(F), y)$ is a scoring rule, but strict consistency of S does not imply strict properness of $\tilde{S}$.

One advantage of constructing distributional forecasts is the availability of the probability integral transform (PIT) as a diagnostic tool, which has been widely used in statistics and econometrics (see for example Dawid, 1984, Diebold, Gunther, and Tay, 1998, Christoffersen, 1998, Gneiting, Balabdaoui, and Raftery, 2007, Mitchell and Wallace, 2011). PIT refers to the basic fact that if a random variable Y has continuous distribution function F then the random variable $U = F(Y)$ has uniform $[0, 1]$ distribution. Thus uniformity of samples of Y suggests that the distribution F has been correctly evaluated. We make extensive use of the PIT in a somewhat different context in Section 5 below.

There is one class of problems in which 'point forecasts' and 'distributional forecasts' coalesce, namely the problem of predicting the success probability in Bernoulli trials, where of course the success probability *is* the distribution. There is an extensive literature on this problem—see Lai et al. (2011) for a recent example where martingale theory is used in a somewhat similar way to Section 6 below.

3.4. **Dynamic Models.** Suppose we observe not just a single variate Y but a sequence $Y_1, Y_2, \ldots$. Any corresponding stochastic model is then a discrete-time process on some probability space $(\Omega, \mathcal{G}, \mathbb{P})$, for which we denote by $F_k(y)$ the conditional distribution of Y_k given $Y_1, \ldots, Y_{k-1}$:

$$F_k(y) = \mathbb{P}[Y_k \leq y | Y_1, \ldots, Y_{k-1}].$$

Suppose that, for some class $\mathcal{F}$ of distributions and for all sequences $y_1, y_2, \ldots$

(i) $F_k(\cdot\,; y_1, \ldots, y_{k-1}) \in \mathcal{F}$;

(ii) For a given statistic $\mathfrak{s}$ there is an identification function V such that for $F \in \mathcal{F}$

$$x \in \mathfrak{s}(F) \Leftrightarrow \mathbb{E}_F[V(x, Y)] = 0.$$

Then when $x_k = \mathfrak{s}(F_k)$ we have

$$\mathbb{E}[V(x_k, Y_k) | Y_1, \ldots, Y_{k-1}] = 0,$$

[6]Gneiting and Raftery (2007) maximise instead of minimising and allow extended real valued scoring rules.

i.e. $B_j \triangleq \sum_{k=1}^{j} V(x_k, Y_k)$ is a martingale. This is the basis of our approach in the remainder of the paper. Under the conditions stated by Steinwart et al. (2014), item (ii) is tantamount to stating that the statistic $\mathfrak{s}$ is elicitable in the class $\mathcal{F}$.

4. Calibration of Predictions

Now let us return to the world of risk management, a dynamic situation in which, at time $k-1$ we have observed a real-valued price series $Y_1, \ldots, Y_{k-1}$ and an $\mathbb{R}^r$-valued series of other data $H_1, \ldots, H_{k-1}$ and wish to make some prediction relating to the behaviour of Y_k.

A *model* for the data is a discrete-time stochastic process $(\tilde{Y}_k, \tilde{H}_k)$ defined on a stochastic basis $(\Omega, \mathcal{G}, (\mathcal{G}_k), \mathbb{P})$. We always take $(\Omega, \mathcal{G}, (\mathcal{G}_k))$ to be the canonical space for an $\mathbb{R}^{1+r}$-valued process, i.e. $\Omega = \prod_{k=1}^{\infty} \mathbb{R}^{1+r}_{(k)}$ (where each $\mathbb{R}^{1+r}_{(k)}$ is a copy of $\mathbb{R}^{1+r}$) equipped with the σ-field $\mathcal{G}$, the product σ-field generated by the Borel σ-field in each factor. For $\omega \in \Omega$ we write $\omega = (\omega_1, \omega_2, \ldots) \equiv ((\tilde{Y}_1(\omega), \tilde{H}_1(\omega)), (\tilde{Y}_2(\omega), \tilde{H}_2(\omega)), \ldots)$. The filtration $(\mathcal{G}_k, k \geq 0)$ is then the natural filtration of the process $(\tilde{Y}_k, \tilde{H}_k)$, with $\mathcal{G}_0 = (\Omega, \emptyset)$. With this set-up, different models amount to different choices of the probability measure $\mathbb{P}$ on the same probability space Ω. Below we will consider families $\mathcal{P}$ of probability measures, and we will use the notation $\mathcal{P} = \{\mathbb{P}^m, m \in \mathfrak{M}\}$, where $\mathfrak{M}$ is an arbitrary index set, to identify different elements $\mathbb{P}^m$ of $\mathcal{P}$. The expectation with respect to $\mathbb{P}^m$ is denoted $\mathbb{E}^m$. It is not assumed that the observed data is a sample function of some model $m \in \mathfrak{M}$, or indeed of any model.

Lemma 4.1. *Let $\mathbb{P}^m$ be any probability measure on $(\Omega, \mathcal{G}, (\mathcal{G}_k))$ as defined above. Then for each $k \geq 1$ there is a regular right-continuous conditional distribution of $\tilde{Y}_k$ given $\mathcal{G}_{k-1}$, i.e. a function $F_k^m : \mathbb{R} \times \Omega \to [0,1]$ such that (i) for a.e. ω, $F_k^m(\cdot, \omega)$ is a distribution function on $\mathbb{R}$ and (ii) for each $x \in \mathbb{R}$,*

$$F_k^m(x, \omega) = \mathbb{P}^m[Y_k \leq x | \mathcal{G}_{k-1}] \quad \text{a.s. } (\mathbb{P}^m).$$

Proof. For $k = 1$, $F_1^m(x, \omega) = \mathbb{P}^m[\tilde{Y}_1 \leq x]$, the unconditional distribution function. For $k > 1$ the assertions of the lemma only involve the finite-dimensional vector r.v.

$$((\tilde{Y}_1, \tilde{H}_1), \ldots, (\tilde{Y}_{k-1}, \tilde{H}_{k-1}), \tilde{Y}_k) \in \mathbb{R}^{k(1+r)-r}.$$

Existence of a regular conditional distribution follows from Theorem 10.2.2 of Dudley (1989). □

With these preliminaries in place, we now want to introduce the concept of *calibration* for a statistic $\mathfrak{s}$ relative to a class of models $\mathcal{P}$. Let $\mathfrak{I}(\mathcal{P})$ denote the set of strictly increasing predictable processes (b_n) on $(\Omega, (\mathcal{G}_k))$ such that $\lim_{n\to\infty} b_n = \infty$ a.s. $\forall \mathbb{P}^m \in \mathcal{P}$; in this context, 'predictable' means that for each k, b_k is $\mathcal{G}_{k-1}$-measurable. Often, b_k will actually be deterministic. A *calibration function* is a measurable function $\ell : \mathbb{R}^2 \to \mathbb{R}$, chosen so that

$$\mathbb{E}^m[\ell(\tilde{Y}_k, \mathfrak{s}(F_k^m)) | \mathcal{G}_{k-1}] = 0 \tag{4.1}$$

for all $\mathbb{P}^m$ in some class $\mathcal{P}$. Formally, this property is equivalent to saying that ℓ is an identification function as defined in Section 3.2, but we use a different notation since here there is an extra ingredient $b \in \mathfrak{I}(\mathcal{P})$, so ℓ is just one component of the pair (ℓ, b). The norming sequence b_n has no direct counterpart in elicitability theory. We will see below that in the case of statistics $\mathfrak{s}$ involving expectations it may be necessary to take random norming sequences and then the conditions for calibration become more complicated.

Definition 4.2. *A statistic $\mathfrak{s}$ is* (ℓ, b)-calibrated *in a set $\mathcal{P} = \{\mathbb{P}^m : m \in \mathfrak{M}\}$ of probability measures on $(\Omega, \mathcal{G})$, where ℓ is a calibration function and $b \in \mathfrak{I}(\mathcal{P})$, if*

$$\lim_{n\to\infty} \frac{1}{b_n} \sum_{k=1}^{n} \ell(\tilde{Y}_k, \mathfrak{s}(F_k^m)) = 0 \quad \mathbb{P}^m\text{−a.s. for each } m \in \mathfrak{M}. \tag{4.2}$$

The criterion (4.2) only depends on realized values of data and numerical values of predictions, in accordance with the weak prequential principle.

In practice we observe the data sequence $(Y_1, H_1, \dots, (Y_{k-1}, H_{k-1})$ and produce an estimate $\pi(k)$, based on some algorithm, for what we claim to be $\mathfrak{s}(F_k)$. The point of the calibration process is to check whether the 'statistics' $\pi(n)$ we produce can reasonably be accepted as relating to some putative 'conditional distribution'. Specifically, the quality of our predictions is gauged by calculating

$$J_n(Y, \pi) = \frac{1}{b_n} \sum_{k=1}^{n} \ell(Y_k, \pi(k)).$$

Calibration is a 'reality check': it says that if $(Y_i, H(i))$ were actually a sample function of some process and we did use the correct predictor $\pi(i) = \mathfrak{s}(F_i)$ then the loss J_n will tend to zero for large n, and this will be true whatever the model generating $Y(i)$, within the class $\mathcal{P}$, so a small value of J_n is evidence that our prediction procedure is well-calibrated. The evidence is strongest when $\mathcal{P}$ is a huge class of distributions and b_n is the slowest-diverging sequence that guarantees convergence in (4.2) for all $\mathbb{P} \in \mathcal{P}$. Calibration is however only a necessary condition. We will see in Section 5.3 below that there can be 'nonsense' predictors that survive the calibration tests while being almost unrelated to the data. This means that, to complete the picture, we need more tests to determine whether our predictions are related to the data according to clearly-stated criteria. An example in the case of VaR estimation will be found in Section 7.1.

5. Quantile forecasting

Quantile forecasting is in a sense the 'dual' of probability forecasting. In the weather forecasting problem described in Section 2.4 the *event* (rain/no rain) is always the same and we forecast the *probabilities* p_n, while in quantile forecasting the probability is fixed, $p_n = 1 - \beta$ where β is the significance level, and the forecaster specifies the *event* (loss $\geq q_n$) by selecting q_n. As in Section 4 our set of models is

$$(\Omega, \mathcal{G}, (\mathcal{G}_k), (\tilde{Y}_k, \tilde{H}_k), \mathbb{P}^m), \quad \mathbb{P}^m \in \mathcal{P}$$

where $\mathcal{P}$ is some class of measures and $F_k^m(x, \omega)$ is the conditional distribution function of $\tilde{Y}_k$ given $\mathcal{G}_{k-1}$ under measure $\mathbb{P}^m \in \mathcal{P}$. Let $\mathfrak{P}$ be the set of all probability measures on $(\Omega, \mathcal{G})$, and define

$$\mathcal{P}_c = \{\mathbb{P}^m \in \mathfrak{P} : \forall k, F_k^m(\cdot\,, \omega) \in \mathcal{F}_c \text{ for almost all } \omega \in \Omega\}. \tag{5.1}$$

Here, $\mathcal{F}_c$ is the set of continuous distribution functions. For risk management applications, the continuity restriction is of no significance; no risk management model would predict positive probability for *specific values* of future prices[7]. So $\mathcal{P}_c$ is the biggest relevant subset of $\mathfrak{P}$.

The following result is a slight extension of a well-known result generally credited to Rosenblatt (1952); as mentioned in Section 3.3 it is widely used in statistics and econometrics, and it is also used by Holzmann and Eulert (2014) in much the same context as here. We give a statement and simple proof below to stress the fact that the result imposes absolutely no restriction on the stochastic basis or the joint distribution of the $\tilde{Y}_k$ beyond the requirement that all conditional distributions be continuous. As before, F_1^m denotes the unconditional distribution function of $\tilde{Y}_1$.

[7]Unless the model is based on Monte-Carlo generated empirical distributions, in which case some form of smoothing would be required.

Proposition 5.1. *Suppose $\mathbb{P}^m \in \mathcal{P}_c$, defined by (5.1) above. Then the random variables $U_k = F_k^m(\tilde{Y}_k)$, $k = 1, 2, \ldots$ are i.i.d. with uniform distribution $U[0,1]$.*

Proof. There are at most countably many intervals $I_k, k = \ldots -1, 0, 1, 2, \ldots$ of positive length such that F_1^m takes constant value v_k on I_k with $v_k < v_{k+1}$. For $x \notin \mathcal{I} = \bigcup_k I_k$, F_1^m is 1-1 and $\mathbb{P}^m[U_1 \le u_1] = \mathbb{P}^m[\tilde{Y}_1 \le (F_1^m)^{-1}(u_1)] = u_1$. Since $\mathcal{I}$ has F_1^m measure 0 we conclude that $U_1 \sim U[0,1]$. Similarly, $U_k \sim U[0,1]$ for each $k > 1$. Now suppose that $U_1, \ldots, U_n$ are independent for some n. Then

$$\begin{aligned}\mathbb{P}^m[U_i \le u_i,\, i = 1, \ldots, n+1] &= \mathbb{E}^m\left[\left(\prod_{i=1}^n \mathbf{1}_{(U_i \le u_i)}\right)\mathbb{P}^m[U_{n+1} \le u_{n+1} | \mathcal{G}_n]\right] \\ &= \mathbb{E}^m\left[\left(\prod_{i=1}^n \mathbf{1}_{(U_i \le u_i)}\right)\right] u_{n+1} = \prod_{i=1}^{n+1} u_i.\end{aligned}$$

Thus all finite-dimensional distributions of (U_i) are i.i.d. $U[0,1]$. □

5.1. **Calibration of quantile estimates.** For $\beta \in (0,1)$ let q_k^m denote the β'th quantile of F_k^m, i.e. $q_k^m = \inf\{x : F_k^m(x) \ge \beta\}$. q_k^m is an $\mathcal{G}_{k-1}$-measurable random variable for each $k > 0$. We use the calibration function ℓ defined at (3.4) above.

Theorem 5.2. *If $\mathbb{P}^m \in \mathcal{P}_c$, defined at (5.1), then for any sequence $b_n \in \mathfrak{I}(\mathcal{P})$,*

$$\frac{1}{b_n}\frac{1}{n^{1/2}(\log\log n)^{1/2}}\sum_{k=1}^n (\mathbf{1}_{(Y_k \le q_k^m)} - \beta) \to 0 \quad \text{a.s. } (\mathbb{P}^n) \tag{5.2}$$

Thus the quantile statistic $\mathfrak{s}(F) = q_\beta$ is (l, b')-calibrated for $\mathcal{P}_c$ in accordance with Definition 4.2, where $\ell(x,q) = \mathbf{1}_{(x \le q)} - \beta$ and $b'_k = b_k(k \log\log k)^{1/2}$.

Proof. By monotonicity of the distribution function, $(Y_k \le q_k^m) \Leftrightarrow (U_k \le F_k^m(q_k^m)) \Leftrightarrow (U_k \le \beta)$. The result now follows from Proposition 5.1 and by applying the Law of the Iterated Logarithm (LIL) (Dudley, 1989, Theorem 12.5.1) to the sequence of random variables $Z_k = \mathbf{1}_{(U_k \le \beta)} - \beta$, which are i.i.d with mean 0 and variance $\beta(1-\beta)$. Indeed, define

$$\zeta(n) = \frac{1}{\sigma(2n\log\log n)^{1/2}}\sum_{k=1}^n Z_k$$

where $\sigma = \sqrt{\beta(1-\beta)}$. Then the LIL asserts that, almost surely,

$$\limsup_{n\to\infty} \zeta(n) = 1, \qquad \liminf_{n\to\infty} \zeta(n) = -1.$$

The convergence in (5.2) follows. □

Of course, if convergence holds in (5.2) then it also holds if we replace the sequence b by b'' such that $b''_n \ge b_n$ for all n. In particular, the conventional relative frequency measure

$$\frac{1}{n}\sum_{k=1}^n (\mathbf{1}_{(Y_k \le q_k^m)} - \beta) \tag{5.3}$$

converges under the same conditions; this also follows directly from the Strong Law of Large Numbers (SLLN) (Dudley, 1989, Theorem 8.3.5.); however, the LIL gives a stronger result.

The striking thing about Theorem 5.2 is that calibration of quantile forecasting is obtained under essentially *no* conditions on the mechanism generating the data. As we shall see below, we cannot expect any such strong result in estimating other risk measures.

Theorem 5.2 is a 'theoretical' result in that (5.2) is a tail property, unaffected by any initial segment of the data. Nonetheless, it is practically relevant to compute the relative frequency

(5.3). As we will show in Section 8 below, doing so can provide convincing evidence that our prediction procedure is well calibrated, i.e. produces the right relative frequency of threshold exceedances, consistent with q_k^m being the true β-quantile of F_k^m. For further evidence, we could examine by statistical test the other claim of Proposition 5.1, namely that the random variables (U_k), and hence the binary variates $\mathbf{1}_{(Y_k \le q_k^m)}$ are independent. We address this issue next.

5.2. **A test for serial dependence.** Given our prediction algorithm and the data return sequence Y_k we generate a sequence $\mathfrak{a} = (a_0, a_1, \ldots)$ of binary r.v. $a_k = \mathbf{1}_{(Y_k \le q_k^m)}$. The above tests give confidence that that $\mathfrak{a}$ is consistent with a model in which $\mathbb{P}[a_k = 1] = \beta$. We now want to test the first 'i' in i.i.d., the null hypothesis being

$$\mathfrak{H}_0 : \text{The } a_k \text{ are i.i.d. with } \mathbb{P}[a_k = 1] = \beta.$$

There are many tests that address this problem; some references were given in Section 3.3. An obvious recourse would be to use a non-parametric test such as the Wald-Wolfowitz 'runs' test (Gibbons and Chakraborti, 2010, §6.2). However, since we already know the marginal probability β, and since it seem seems unlikely that a_j and a_k will fail to be essentially independent when $j \ll k$, it seems appropriate to use a test for 'local' dependence. For this, a possible set of alternatives is[8]

$$\mathfrak{H}_{\beta,\theta} : \begin{array}{l} \mathfrak{a} \text{ is a sample from a 2-state Markov chain} \\ \text{with stationary distribution } \mathbb{P}[a_k = 1] = \beta. \end{array}$$

Under $\mathfrak{H}_{\beta,\theta}$ the transition probabilities are

$$\begin{aligned} \mathbb{P}[a_0 = 1] &= \beta \\ \mathbb{P}[a_k = 1 | a_{k-1} = 0] &= \theta \\ \mathbb{P}[a_k = 1 | a_{k-1} = 1] &= \theta'. \end{aligned}$$

The stationary distribution is β if

$$\begin{aligned} \beta = \mathbb{P}[a_1 = 1] &= \mathbb{P}[a_1 = 1 | a_0 = 0](1-\beta) + \mathbb{P}[a_1 = 1 | a_0 = 1]\beta \\ &= \theta(1-\beta) + \theta'\beta. \end{aligned}$$

Thus θ and θ' are related, for given β, by

$$\theta' = 1 - \frac{1-\beta}{\beta}\theta, \tag{5.4}$$

so $\mathfrak{H}_{\beta,\theta}$ is a 1-parameter family indexed by $\theta \in [0,1]$, when $\beta \ge \frac{1}{2}$. Assuming $\beta \ge \frac{1}{2}$ is no loss of generality since otherwise we can interchange the roles of '0' and '1'. The i.i.d. case is $\theta = \theta' = \beta$. The log likelihood ratio $\mathrm{LLR}_\theta^{\mathrm{n}}(\mathfrak{a}) = d\mathbb{P}_{\beta,\theta}/d\mathbb{P}_0$ is given by

$$\mathrm{LLR}_\theta^{\mathrm{n}}(\mathfrak{a}) = \text{const} + n_1 \log(1-\theta) + n_2 \log(1-\theta f) + (n - n_1 - n_2)\log(\theta),$$

where $f = (1-\beta)/\beta$ and n_1, n_2 are the numbers of 00, 11 pairs respectively in $\mathfrak{a}$. We denote $\bar{n}_i = n_i/n$, $i = 1, 2$.

Theorem 5.3. *Suppose $\beta \ge \frac{1}{2}$. Then*

(i) The maximum likelihood estimate of θ is

$$\hat{\theta}_\beta(\mathfrak{a}) = \frac{1}{2f}\left(1 - \bar{n}_2 + f(1-\bar{n}_1) - \sqrt{(f-c_1)^2 + 4f(c_1-c_2)}\right) \tag{5.5}$$

where $c_1 = 1 - f\bar{n}_1 - \bar{n}_2$, $c_2 = 1 - \bar{n}_1 - \bar{n}_2$.

[8]Christoffersen (1998) considers Markov chain alternatives but without the stationarity condition.

(ii) The estimator is consistent: under $\mathfrak{H}_{\beta,\theta}$, almost surely as $n \to \infty$

$$\begin{aligned} \bar{n}_1 \to n_1^* &= (1-\theta)(1-\beta) \\ \bar{n}_2 \to n_2^* &= \beta - (1-\beta)\theta, \end{aligned}$$

and $\hat{\theta}_\beta(n_1^*, n_2^*) = \theta$.

The proof of this result is given in Propositions A.1 and A.2 in Appendix A below.

The hypothesis $\mathfrak{H}_0$ that the a_i are independent is equivalent to $\mathfrak{H}_{\beta,\theta}$ with $\theta = \beta$. In this case, $n_1^* = (1-\beta)^2, n_2^* = \beta^2$ and $\hat{\theta}_\beta(n_1^*, n_2^*) = \beta$. We can use the results of Theorem 5.3 to define a 2-sided test at significance level γ in which $\mathfrak{H}_0$ is rejected if $\hat{\theta}(\bar{n}_1, \bar{n}_2) \notin [t_1, t_2]$ where the intervals $[0, t_1)$ and $(t_2, 1]$ each have probability $\frac{\gamma}{2}$ under $\mathfrak{H}_0$. The endpoints t_1, t_2 are easily determined by simulation. Tables 5.1 and 5.2 give their values when $\beta = 0.9, 0.95$ respectively, for four values of the significance level γ.

An application of this test is given in Section 8 below.

γ	Data length 250		Data length 500		Data length 1000	
1%	0.7038	1.0000	0.7785	1.0000	0.8201	0.9672
5%	0.7676	1.0000	0.8103	0.9758	0.8418	0.9538
10%	0.7926	1.0000	0.8272	0.9652	0.8519	0.9450
50%	0.8643	0.9437	0.8728	0.9281	0.8823	0.9200

TABLE 5.1. Confidence intervals t_1, t_2 for estimator $\hat{\theta}_\beta$, $\beta = 0.90$.

γ	Data length 250		Data length 500		Data length 1000	
1%	0.6080	1.0000	0.7854	1.0000	0.8516	1.0000
5%	0.7600	1.0000	0.8398	1.0000	0.8800	1.0000
10%	0.8012	1.0000	0.8648	1.0000	0.8940	1.0000
50%	0.9133	1.0000	0.9249	1.0000	0.9308	0.9732

TABLE 5.2. Confidence intervals t_1, t_2 for estimator $\hat{\theta}_\beta$, $\beta = 0.95$.

5.3. **A 'nonsense' quantile predictor.** Even if a quantile predictor passes both the calibration and independence tests it may still be seriously deficient. A striking example illustrating this for the 95% quantile was given by Holzmann and Eulert (2014, §3.1). The quantile predictor is set at some very high level h on 95 out of every 100 dates, and at a very low level l on the remaining 5 dates. Then the empirical exceedence frequency will be almost exactly 5%, although the predictor is barely related to the data. A variant of this example, already given by Engle and Manganelli (2004), would be to take an i.i.d. Bernoulli sequence B_k with $\mathbb{P}[B_k = 1] = 0.05$ and define our quantile predictor as $\hat{q}_k = lB_k + h(1 - B_k)$. Then almost always $\mathbf{1}_{Y_k \geq \hat{q}_k} = B_k$, so this predictor will pass both the LIL test of Theorem 5.2 and the independence test of Section 5.2.

Examples of this sort pose a major challenge to verification of prediction. In the case of distributional prediction, as discussed in Section 3.3, Gneiting et al. (2007) proposed a diagnostic approach based on 'maximising the sharpness of the predicted distributions subject to calibration'. The idea of 'sharpness' is that, given two distributions, the one with the minimum dispersion (as measured by an inter-quantile range, for example) should be preferred. While this may be appropriate in certain applications such as predicting macroeconomic variables, it is not without controversy, see Mitchell and Wallace (2011). In any case, the principle is not

applicable for point forecasts as considered in this paper. Reacting to the 'nonsense' example, a reasonable criterion might be that predictor A is preferred to predictor B, given that both survive calibration tests, if A is more sensitive to the data than B. If prediction is based on a data vector $X = (X_1, \dots, X_n)$, where X_n is the most recent data point, we could compute directional derivatives $\partial_Z A = \lim_{\epsilon \downarrow 0}(A(X+\epsilon Z) - A(X))/\epsilon$ and $\partial_Z B$ for a range of deterministic pertubation vectors Z and prefer A to B if A has the greater average derivative. Obvious choices of Z could be $\mathbf{1} = (1, \dots, 1)$ or $Z_k = \alpha^{n-k}$ with $\alpha \in (0, 1)$ if sensitivity to recent data is thought to be more important. With $Z = \mathbf{1}$ the nonsense predictor has sensitivity practically equal to zero, while the quantile predictor introduced in Section 8 has sensitivity close to 1. Further investigation of these ideas is a subject for future research. A completely different diagnostic that successfully separates these two predictors is described in Section 7.1 below. One thing all investigators are agreed upon is that, while calibration is—in accordance with the weak prequential principle—a property of the data and predictors jointly, diagnostics beyond calibration are functions of the predictors *only*.

6. Risk Measures Involving Mean Values

Risk measures such as CVaR involve integration with respect to the conditional distribution functions F_k^m. In this section we will consider the straight prediction problem of estimating the conditional means

$$\mu_k^m = \int_{\mathbb{R}} x F_k^m(dx). \tag{6.1}$$

We must assume that the class of candidate models is at most

$$\mathcal{P}^1 = \left\{ \mathbb{P}^m \in \mathfrak{P} : \forall k, \int_{\mathbb{R}} |x| F_k^m(dx) < \infty \right\}.$$

In this context, continuity of the conditional distributions is not required, so $\mathcal{P}^1$ is not a subset of $\mathcal{P}_c$. In fact, this problem is general enough to include risk measures of the form $\int f(x) F_k^m(dx)$ for general functions f: we can simply define a new model class $(\tilde{Y}', \tilde{H}')$ where $\tilde{Y}_k' = f(Y_k)$ and $\tilde{H}_k' = (Y_k, H_k)$. Note that if f is an option-like function such as $f(x) = (x-K)^+$ then $f(\tilde{Y}_k) = 0$ with positive probability for some measures $\mathbb{P}^m$, so it is convenient that we do not require $\mathbb{P}^m \in \mathcal{P}_c$.

6.1. Universality.

The first question to ask is whether we can get any 'universal' result, similar to Theorem 5.2, for estimating μ_k^m, by using the i.i.d. sequence U_k of Proposition 5.1. The answer appears to be no. What makes Theorem 5.2 work is the equality

$$\mathbf{1}_{(\tilde{Y}_k \le q_k^n)} - \beta = \mathbf{1}_{(U_k \le \beta)} - \beta,$$

so by transforming the variables we obtain the universal calibration function $l(u, \beta) = \mathbf{1}_{(u \le \beta)} - \beta$. In the case of expected value prediction the natural criterion is

$$\frac{1}{n} \sum_{k=1}^{n} (\tilde{Y}_k - \mu_k^m) \to 0.$$

Mapping the two variables in the kth term through the distribution function F_k^m gives us a summand

$$U_k - F_k^m(\mu_k^m).$$

This translates into a universal calibration function if and only if there is a constant c such that

$$F_k^m(\mu_k^m) = c \quad \text{a.s. for all } \mathbb{P}^m, \tag{6.2}$$

meaning that μ_k^m coincides with a *fixed quantile* c of F_k^m. But if that is the case the problem reduces to quantile estimation and the results of Section 5 apply. The only natural example of this is the situation where each distribution function F_k^m is symmetric around its mean value, when (6.2) holds with $c = \frac{1}{2}$. This is not a relevant class from the risk-management perspective, but see Abdous and Remillard (1995) where the relations between quantiles and expectiles are examined in greater detail.

6.2. **Martingale analysis.** To proceed further, we need to make use of martingale properties. If we define

$$X_k = \tilde{Y}_k - \mu_k^n, \qquad S_n = \sum_{k=1}^{n} X_k \tag{6.3}$$

with $S_0 = 0$, then S_n is a zero-mean $\mathbb{P}^m$-martingale since $\mathbb{E}^m[X_k|\mathcal{G}_{k-1}] = 0$. We want to determine calibration conditions by using the SLLN for martingales. In this subject, a key role is played by the *Kronecker Lemma* of real analysis.

Lemma 6.1. *Let x_n, b_n be sequences of numbers such that $b_n > 0$, $b_n \uparrow \infty$, and let $u_n = \sum_{k=1}^{n} x_n/b_n$. If $u_n \to u_\infty$ for some finite u_∞ then*

$$\lim_{n\to\infty} \frac{1}{b_n} \sum_{k=1}^{n} x_k = 0.$$

The *martingale convergence theorem* states that if S_n is a zero-mean martingale on a filtered probability space and there is a constant K such that $\mathbb{E}|S(n)| \leq K$ for all n, then $S_n \to S(\infty)$ a.s. where $S(\infty)$ is a random variable such that $\mathbb{E}|S_\infty| < \infty$.

Now let X_k, S_k be as defined at (6.3) above, and let Z_k be a *predictable* process, i.e. Z_k is $\mathcal{G}_{k-1}$-measurable, such that $Z_k > 0$ and $Z_k \uparrow \infty$ a.s. Let $X_k^Z = X_k/Z_k$ and $S^Z(n) = \sum_1^n X_k^Z$. Then S_n^Z is a martingale[9], since

$$\mathbb{E}^m[X_k^Z|\mathcal{G}_{k-1}] = \frac{1}{Z_k}\mathbb{E}^m[X_k|\mathcal{G}_{k-1}] = 0.$$

If we can find Z_k such that $\mathbb{E}^m|S^Z(n)| < c_Z$ for some constant c_Z then S^Z converges a.s. and hence by the Kronecker lemma

$$\frac{1}{Z(n)}S(n) = \frac{1}{Z(n)}\sum_{k=1}^{n}(\tilde{Y}_k - \mu_k^n) \to 0 \quad \text{a.s.}$$

We have shown

Proposition 6.2. *Under the above conditions, the statistic $\mathfrak{s}(F) = \int xF(dx)$ is (ℓ, Z) calibrated in the class $\mathcal{P}^1$, according to the Definition* (4.2)*, where $\ell(x, \mu) = x - \mu$.*

Note that the calibration function ℓ is the one derived from elicitability, see (3.2). The above proposition is of course useless as it stands, because no systematic way to specify the norming process Z_k has been provided. We can partially resolve this problem by moving to a setting of *square-integrable martingales* (see Williams, 1991, Chapter 12). If $S(n) \in L_2$ we define the 'angle-brackets' process $\langle S\rangle_n$ by

$$\langle S\rangle_n = \sum_{k=1}^{n} \mathbb{E}[X_k^2|\mathcal{G}_{k-1}]. \tag{6.4}$$

This is the increasing process component in the Doob decomposition of the submartingale $S^2(n)$.

[9] S^Z is a stochastic integral, or 'martingale transform' of S.

Proposition 6.3 (Williams 1991)**.** *If $S(n)$ is a square-integrable martingale then $S(n)/\langle S\rangle_n \to 0$ on the set $\{\omega : \langle S\rangle_\infty(\omega) = \infty\}$.*

Proof Define the martingale $W(n) = \sum_{k\le n} X_k/(1+\langle S\rangle_k)$, for which

$$\begin{aligned}\mathbb{E}[(W(n)-W(n-1))^2|\mathcal{G}_{n-1}] &= \frac{1}{(1+\langle S\rangle_n)^2}(\langle S\rangle_n - \langle S\rangle_{n-1}) \\ &\le \frac{1}{1+\langle S\rangle_{n-1}} - \frac{1}{1+\langle S\rangle_n} \quad \text{a.s.}\end{aligned}$$

It follows that $\langle W\rangle_\infty \le 1$. From Williams (1991, Theorem 12.13) this implies that $\lim_n W_n$ exists, and hence from the Kronecker lemma that $S(n)/\langle S\rangle_n \to 0$ as long as $\langle S\rangle_n \uparrow \infty$. □

Proposition 6.3 shows that in the square-integrable case we can take $Z = \langle S\rangle$ in Proposition 6.2. However, we cannot use $\langle S\rangle$ as it stands because it does not satisfy the weak prequential principle, which requires that the norming sequence be calculable using only observed data and numerical values of estimates. To achieve this, we follow a line of reasoning pursued by Hall and Heyde (1980), relating the predictable quadratic variation $\langle S\rangle_n$ to the realized quadratic variation

$$Q_n = \sum_{k=1}^n (S_k - S_{k-1})^2 = \sum_{k=1}^n Y_k^2.$$

As Hall and Heyde point out, the two random variables Q_n and $\langle S\rangle_n$, defined at (6.4), have the same expectation, and we are interested in the ratio $Q_n/\langle S\rangle_n$. To get the picture, consider the case where the Y_k are i.i.d. with variance σ^2. Then $\langle S\rangle_n = \sigma^2 n$ and

$$\lim_{n\to\infty} \frac{Q_n}{\langle S\rangle_n} = \frac{1}{\sigma^2} \lim_{n\to\infty} \frac{1}{n}\sum_{k=1}^n Y_k^2 = 1 \quad \text{a.s.} \tag{6.5}$$

by the SLLN. In the general, martingale, case we may or may not have convergence as in (6.5), as Hall and Heyde (1980) show. We will not present details of their analysis here but content ourselves with the following definition.

Definition 6.4. *Let $\mathcal{P}^e \subset \mathfrak{P}$ be the set of probability measures $\mathbb{P}^m$ such that*
(i) $\forall k,\ \tilde{Y}_k \in L_2(\mathbb{P}^m)$.
(ii) $\lim_{n\to\infty}\langle S\rangle_n = \infty$ a.s. $\mathbb{P}^m$, where S_n is defined at (6.3)*.*
(iii) There exists $\epsilon_m > 0$ such that $Q_n/\langle S\rangle_n > \epsilon_m$ for large n, a.s. $\mathbb{P}^m$.

We can now state our final result.

Theorem 6.5. *The mean statistic $\mathfrak{s}(F) = \int xF(dx)$ is (ℓ, Q_n) calibrated for the class $\mathcal{P}^e$, where*

$$\ell(x,\mu) = x - \mu.$$

Proof. Suppose $\mathbb{P}^m \in \mathcal{P}^e$. Conditions (i) and (ii) of Definition 6.4 imply that $S(n)/\langle S\rangle_n \to 0$ by Proposition 6.3. Using condition (iii) we have

$$\left|\frac{S(n)}{Q_n}\right| = \frac{\langle S\rangle_n}{Q_n}\left|\frac{S(n)}{\langle S\rangle_n}\right| \le \frac{1}{\epsilon_m}\left|\frac{S(n)}{\langle S\rangle_n}\right| \quad \text{for large } n.$$

The result follows. □

As we see, significant conditions must be imposed to secure consistency of mean-type estimates, in contrast to the situation for quantile estimates (Theorem 5.2) where almost *no* conditions are imposed. Theorem 5.2 is a LIL-based result whereas Theorem 6.5 is based on the SLLN. There is a sizable literature on the LIL for martingales (see Hall and Heyde, 1980, again),

but a number of quite intricate conditions are required, none of which would be checkable in the context of mean estimation, so it does not seem worth pursuing this avenue here.

Arguments based on martingale convergence have been used in similar contexts by Dawid and Vovk (1999) and Lai et al. (2011).

Verifying the validity of mean-based estimates is always more problematic than the same problem for quantile-based statistics. In fact the whole process of mean estimation is more problematic because, just from the basic definition (6.1), the mean depends in an essential way on the tail of the distribution function F and, in any situation involving real data rather than model-generated data, we run out of data at some point in trying to estimate the tail, but the unestimated part may contribute significantly to the mean. We discuss in the next section question of CVaR estimation where this difficulty can be seen very clearly.

7. Estimating CVaR

This section focuses on computation of CVaR. It has been pointed out by Cont, Deguest, and Scandolo (2010) that CVaR is excessively sensitive to small changes in the data sequence from which it is computed. Here we wish to make the more general point that any mean calculation depends on the tails of the distribution in ways that cannot be easily controlled. This is discussed in Section 7.2 below but first, in Section 7.1 we introduce the characterisation, due to Rockafellar and Uryasev (2000), of CVaR as the minimum of a certain convex function. Recent work by Holzmann and Eulert (2014) has highlighted the relation between this result and elicitability.

7.1. CVaR as the solution to a minimisation problem.

Let F belong to the set $\mathcal{F}_{c\uparrow}$ of continuous and strictly increasing distribution functions on $\mathbb{R}^+$. From (1.1), the CVaR at level β can be expressed as

$$\mathrm{CVaR}_\beta(F) = q_\beta + \frac{1}{1-\beta}\int_{q_\beta}^{\infty}(y-q_\beta)F(dy) = \frac{1}{1-\beta}\int_\beta^1 q_\tau d\tau.$$

where q_τ is the unique τ-quantile of F. Rockafellar and Uryasev (2000, 2002) give a charaterisation of CVaR as the solution to a minimization problem given as follows.

Proposition 7.1. *For $x \in \mathbb{R}, \beta \in (0,1)$ let*

$$\Psi_\beta^F(x) = x + \frac{1}{1-\beta}\int_x^\infty (y-x)F(dx).$$

Then

$$\mathrm{CVaR}_\beta(F) = \min_{x\in\mathbb{R}} \Psi_\beta^F(x) = \Psi_\beta^F(q_\beta).$$

It has been noted by Holzmann and Eulert (2014, §3.2) that this result is closely related to the elicitability properties of the quantile q_β. Recall from Section 3.1.2 that score functions for the quantile take the form $S(x,y) = (\mathbf{1}_{x\geq y} - \beta)(g(x) - g(y))$. If we take $g(x) = x/(1-\beta)$ then we find that

$$S(x,y) = x + \frac{1}{1-\beta}(y-x)\mathbf{1}_{y>x} - \frac{y}{1-\beta} \;\stackrel{\Delta}{=}\; S^*(x,y) - \frac{y}{1-\beta}. \tag{7.1}$$

The term $y/(1-\beta)$ plays no role in the minimization, and $\mathbb{E}_F[S^*(x,Y)] = \Psi_\beta^F(x)$. Thus the minimum value of $\mathbb{E}_F[S^*(x,Y)]$, achieved at $x = q_\beta$, is exactly CVaR_β.

This result gives us a diagnostic test for comparing VaR estimators that have survived calibration and independence tests. Given a data sequence $Y_1, Y_2, \ldots$, let $\hat{q}_k^m, k = 1, 2\ldots$ be the

sequences of β-quantile predictions produced by two algorithms $m = 1, 2$. Then we prefer predictor 1 to predictor 2 if

$$\frac{1}{n}\sum_{k=1}^{n} S^*(\hat{q}_k^1, Y_k) < \frac{1}{n}\sum_{k=1}^{n} S^*(\hat{q}_k^2, Y_k). \tag{7.2}$$

This procedure of comparing realized average scores is known as a Diebolt-Mariano test (Diebold and Mariano, 1995), see also Gneiting (2011, §1.1). There is no rigorous justification for the test without extra mixing or other assumptions, but nonetheless it proves to be effective. As an example, let $\hat{q}_k^1$ be the predictors of 90% return quantiles for the FTSE100 data of Figure 2.1 produced by the algorithm (8.1), (8.2) described in Section 8 below, and let $\hat{q}_k^2$ be Holzmann and Eulert's 'nonsense' predictor as described in Section 5.3 above. The two parameters l, h for this algorithm are taken as $l = -0.06, h = +0.06$, which are close to being lower and upper bounds for the return sequence, see Figure 2.1(b). We compute the averages as in (7.2) fixing $n = 500$ but taking a moving window of data. Specifically, we compute

$$x_j^m = \frac{1}{500}\sum_{k=j}^{j+499} S^*(\hat{q}_k^m, Y_k), \quad j = 1, \dots, 1000,\ m = 1, 2. \tag{7.3}$$

As can be seen from Figure 7.1, the algorithm of Section 8 is consistently and decisively preferred to the nonsense algorithm.

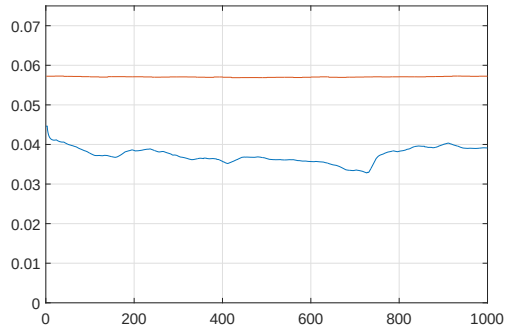

FIGURE 7.1. Plot of x_j^1 and x_j^2, defined by (7.3), against $j = 1, \dots, 1000$. Lower curve is x_j^1.

7.2. **Basic problems in CVaR estimation.** In Section 2.2 we saw that the empirical distribution of returns for the FTSE100 data set displayed power tails (defined precisely in the footnote there) with tail index 2.35 on the left (=loss) side. We should not read too much into this since it is not claimed that the returns are samples from the same distribution, but nevertheless it does add credibility to the idea of considering power-tail distributions as candidates for a model in the sense defined in Section 4.

To clarify the difficulty in CVaR estimation, consider the following proposition, in which F is supposed to have *exact* power tail. Its proof is a simple computation.

Proposition 7.2. *Let $0 < \beta < \eta < 1$ and F be a continuous distribution function on $\mathbb{R}^+$ such that for $x \geq q_\eta^+$*

$$F(x) = 1 - (1-\eta)\left(\frac{x}{q_\eta}\right)^{-\kappa}$$

where $\kappa > 1$. *Then*

$$\mathrm{CVaR}_\beta(F) = \frac{1}{1-\beta}\left(\int_\beta^\eta q_\tau d\tau + \frac{\kappa}{\kappa-1}(1-\eta)q_\eta\right). \tag{7.4}$$

It will be seen in the next section that quantile estimation for financial data is something that can be achieved convincingly for significance levels out to 95% at least. A further point is that the representation clearly relates to the idea of estimating distributions by estimating a series of quantiles, see Cervera and Muñoz (1996) or Gneiting and Raftery (2007, §6.1) for further details. Suppose we wish to compute CVaR_β and can reliably estimate quantiles q_τ for $\tau \le \eta$ but not beyond η where the data has dried up. Then the first term on the right of (7.4) and the value of q_η are known, but the result also depends on the value of κ, and $\mathrm{CVaR}_\beta(F) \to +\infty$ as $\kappa \downarrow 1$. To place an upper bound on CVaR requires a reliable estimate for the tail index κ but by definition this is impossible to obtain. The conclusion is that any estimate of CVaR depends on *a priori* assumptions about tail behaviour that cannot be verified on the basis of any finite data set, however large.

Various expedients can be employed:

(i) If the empirical return data exhibits power tails, for example the FTSE100 data where the left (=loss) tail index is $\kappa = 2.35$, then use this value beyond the last point where the quantiles can be accurately estimated.

(ii) Use an econometric model. Any model implies tail behaviour, which might be time and data-dependent.

(iii) Use methods based on extreme-value theory (Embrechts et al., 1997).

(iv) Extrapolation: given reliable estimates for q_β and q_η and assuming one is already in the tail regime at q_β one can back out the implied value of κ. This is however likely to be a very noisy estimate.

(v) Cont et al. (2010) suggest modifying the definition of CVaR_β to

$$\frac{1}{\eta-\beta}\int_\beta^\eta q_\tau d\tau, \quad \text{for some } \eta < 1,$$

providing a robustly computable statistic.

(vi) Kou et al. (2013) propose replacing CVaR by CMVaR, the conditional *median* loss beyond VaR. Clearly, $\mathrm{CMVaR}_\beta = \mathrm{VaR}_\eta$ with $\eta = (1+\beta)/2$, so computation reduces to VaR estimation.

All of these have their disadvantages. Item (i) assumes a relationship between the empirical distributions and the conditional distributions of CVaR which cannot be rigorously justified. Item (ii) is a broad-brush approach which typically uses empirical tail estimates to infer a suitable choice of i.i.d. drivers for the model. In item (iii), extreme-value theory is an analysis of i.i.d. samples and in the present context is best suited to the very high significance levels β of external risk management. Item (iv) is predicated on power tail and is certainly not 'robust', while the remaining two cut off the tail completely at some point, which might miss real risk. Both of these require reliable VaR estimates up to level $\eta > \beta$ where β is the level at which CVaR is required.

From a practical perspective the purpose of computing CVaR is to establish some threshold beyond VaR such that the gap (CVaR − VaR) provides an adequate cushion against extreme losses. From that point of view, CMVaR seems the right choice in that it is easy to compute, has a clear statistical meaning and has axiomatic support (Kou et al., 2013).

8. An algorithm for quantile forecasting

By way of illustration, we present in this section a very simple data-driven algorithm for producing 1-week ahead forecasts of the upper 10% or 5% quantiles of returns on the FTSE100 index. The data is weekly values of the index 1994-2013. Figure 2.1 shows the index values and the series of weekly returns. This series exhibits the usual 'stylized facts': gross variations in volatility and apparent non-stationarity.

To predict quantiles of the returns the traditional approach is to estimate the parameters of an econometric model such as GARCH or EGARCH and then compute the 1-week ahead conditional distribution. However, a much simpler data-driven approach seems competitive in terms of calibration. As always, we are only checking necessary conditions with these tests. We start with the 90% quantile. As a first step we compute, at time step k, an empirical 90% quantile of the most recent 20 values $r_{k-19}, \dots, r_k$. The largest such quantile is of course just the 2nd largest of the 20 values, and this is our predicted quantile $\hat{q}_{k+1}$ for r_{k+1}. Perfect calibration would mean that on average the realized value exceeds the predicted quantile 10% of the time. Figure 8.1 shows the achieved calibration, i.e. graphs

$$y_k = \frac{1}{k}\sum_{j=1}^{k} \mathbf{1}_{(r_j > \hat{q}_j)}.$$

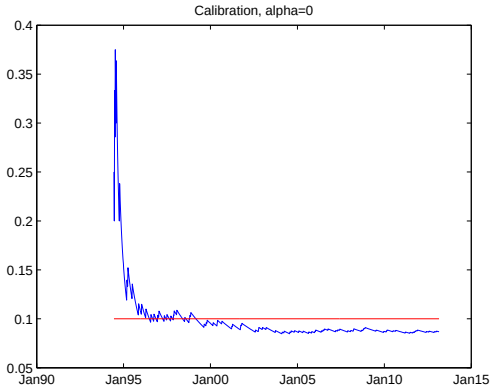

Figure 8.1. FTSE100 Calibration of initial algorithm.

As can be seen, the algorithm is slightly miscalibrated in that the average exceedance rate is around 9% rather than the desired 10%, showing that the average threshold is too high. This may be related to the fact that we chose the largest possible quantile, but in any case is easily rectified by a simple feedback or *adaptive* mechanism that corrects for errors in the achieved performance. Specifically, the new quantile prediction is

$$\check{q}_{k+1} = \hat{q}_{k+1} + \varphi(\check{y}_k - 0.1) \tag{8.1}$$

where

$$\check{y}_k = \frac{1}{k}\sum_{j=1}^{k} \mathbf{1}_{(r_j > \check{q}_j)} \tag{8.2}$$

and φ is a parameter. The performance of this algorithm and the original one are shown in Figure 8.2(a), while Figure 8.2(b) shows the sequence of thresholds produced by the algorithm, which vary dramatically over time. (The straight line is the median threshold, around 0.028.)

The value of φ chosen was 1.2. Performance is not very sensitive to this value, but some number greater than one speeds up convergence of the algorithm.

k	Min %	Max %	Bernoulli SD
50	8.96	11.22	4.24%
100	9.09	11.22	3.00%
250	9.50	10.53	1.90%
500	9.67	10.33	1.84%

TABLE 8.1. Calibration performance of adaptive algorithm

Table 8.1 quantifies the performance of the algorithm. Each row of the table shows, for the number of weeks k in the left-hand column, the minimum and maximum exception frequencies $\min\{\check{y}_k, \ldots, \check{y}_{1500}\}$ and $\max\{\check{y}_k, \ldots, \check{y}_{1500}\}$, with $\check{y}_j$ given by (8.2), experienced beyond that time to the end of the sample. For comparison, the last column shows the standard deviation $\sqrt{p(1-p)/k}$ of the average of k independent Bernoulli trials with success probability $p = 0.1$. All numbers are expressed in percentage terms. It appears that the deviations from the theoretical 10% exception frequency are well within the sampling error expected under the i.i.d. hypothesis. Stable calibration is maintained right through the financial crisis period of 2007-09, though of course the actual thresholds fluctuate widely in response to market conditions.

Kuester, Mittnik, and Paolella (2006) present a highly informative and extremely thorough comparative study of VaR estimation techniques. They summarize their conclusions as follows:

> In this study we compare the out-of-sample performance of existing methods and some new models for predicting value-at-risk (VaR) in a univariate context. Using more than 30 years of the daily return data on the NASDAQ Composite Index, we find that most approaches perform inadequately, although several models are acceptable under current regulatory assessment rules for model adequacy. A hybrid method, combining a heavy-tailed [..] GARCH filter with an extreme value theory-based approach, performs best overall.

Given this conclusion and the quite high computational demands of the methods surveyed, it does appear from the results presented here that data-driven methods, including techniques such as reinforcement learning (Dempster and Leemans, 2006) merit further investigation.

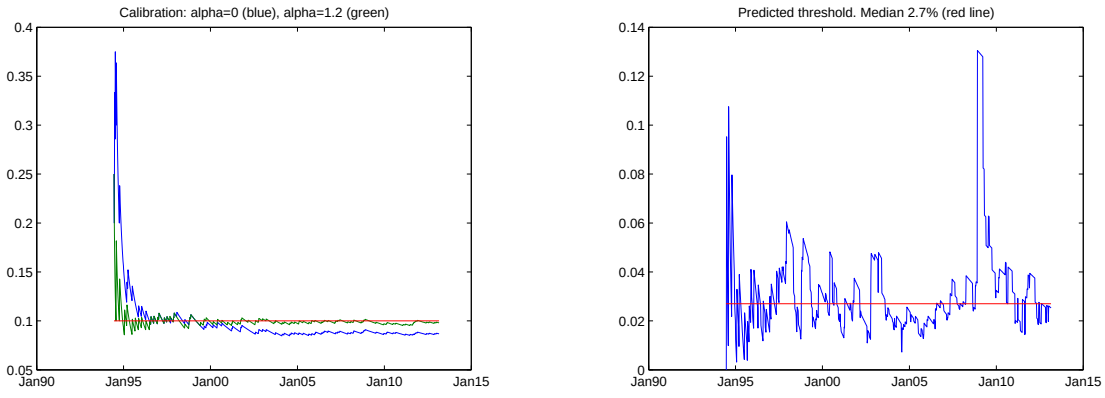

(a) Calibration (b) Quantile estimates

FIGURE 8.2. Performance of feedback algorithm

8.1. **Testing for serial dependence.** We now implement the test for serial dependence introduced in Section 5.2. The null hypothesis is

$$\mathfrak{H}_0 : \text{The } Y_k \text{ are i.i.d. with } \mathbb{P}[Y_k = 1] = \mu.$$

$\mathfrak{H}_0$ is rejected if $\hat{q}_\beta \notin [t_1, t_2]$; these intervals are specified for various significance levels and data lengths in Tables 5.1 and 5.2.

First, we take $\beta = 0.9$ as above and run the test on the maximum data length 1500 for the FTSE100 return series. The calibration results, shown in Figure 8.3(a) are similar to the shorter data length results of Figure 8.2. For the independence test we determine the relative frequencies $\bar{n_1}, \bar{n_2}$ and then calculate $\hat{q}(\bar{n_1}, \bar{n_2})$. From Theorem 5.3, the limiting values of $\bar{n_1}, \bar{n_2}$ under $\mathfrak{H}_0$ are $(1-\beta)^2$, $\beta^2 = 0.01, 0.81$ respectively. The values obtained in our test were

$$\begin{aligned} \bar{n}_1(1500) &= 0.0100 \\ \bar{n}_2(1500) &= 0.8120 \\ \hat{q}_{0.9}(\bar{n}_1, \bar{n}_2) &= 0.8980. \end{aligned}$$

The agreement with the theoretical limiting values is almost perfect and the value of $\hat{q}_{0.9}$ is within 20 basis points of the correct value under independence. This test is however based on the entire 30-year data sequence and gives us only one estimate. A better evaluation is to take running estimates over shorter periods. Figure 8.3(b) shows the results of estimates with a moving window of length 500. This is a plot of z_k against k where

$$z_k = \frac{1}{500} \sum_{j=k-501}^{k} \mathbf{1}_{(r_j > \check{q}_j)} \tag{8.3}$$

The confidence intervals in Table 5.1 show that $\mathfrak{H}_0$ would only occasionally be rejected at the 5% significance level.

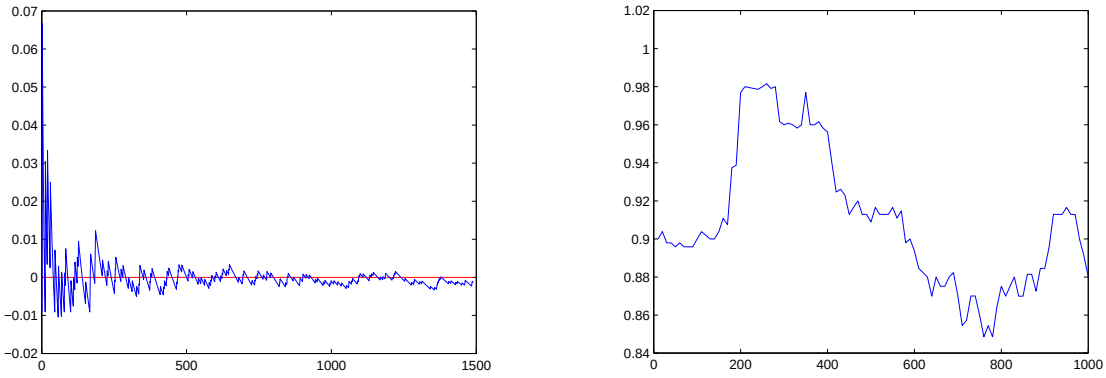

(a) Whole-sample.

(b) Running calibration z_k of (8.3) with data window 500.

FIGURE 8.3. Calibration with $\beta = 0.90$.

Finally, we repeated these tests at the industry standard value $\beta = 0.95$. The results are shown in Figure 8.4. The quantile prediction algorithm is the same as before except that we now take the largest, rather than the 2nd largest, of the previous 20 returns as our predictor. Calibration, shown in Figure 8.4(a), is only slightly less satisfactory than before. Turning to the independence test, the limiting values of $\bar{n_1}, \bar{n_2}$ in this case are $(1-\beta)^2$, $\beta^2 = 0.0025, 0.9025$,

while the achieved values were

$$\begin{aligned} \bar{n}_1(1500) &= 0.0027 \\ \bar{n}_2(1500) &= 0.9007 \\ \hat{q}_{0.95}(\bar{n}_1, \bar{n}_2) &= 0.9481. \end{aligned}$$

The value of $\hat{q}_{0.95}$ is still within 20 bp of the theoretical 'independence' value.

In the running 500 test shown in Figure 8.4(b) the estimates never fall outside the range reported in Table 5.2, even at the 50% significance level. However this test is somewhat less satisfactory in that the upper barrier t_2 is always equal to 1 in this case, so the test reduces to a 1-sided one.

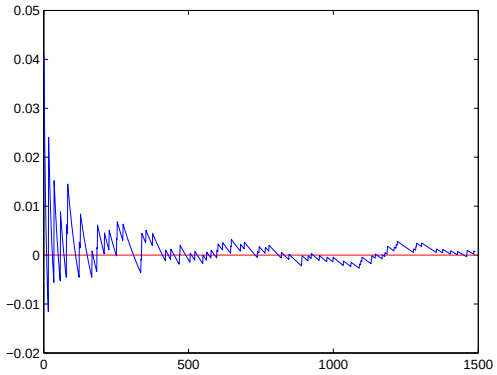

(a) Whole-sample calibration.

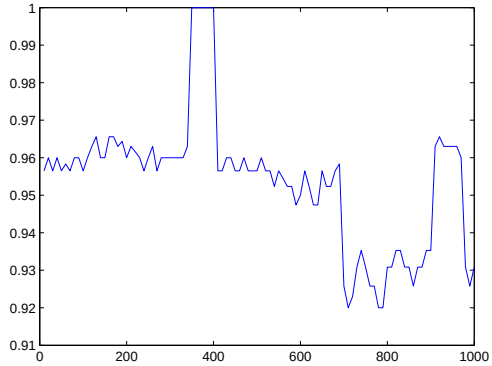

(b) Running calibration z_k of (8.3) with data window 500.

FIGURE 8.4. Calibration with $\beta = 0.95$.

APPENDIX A. THE MARKOV CHAIN MODEL

In model $\mathfrak{H}_{\mu,\theta}$ the transition probabilities of the chain are as shown in Table A.1, where $f = (1-\mu)/\mu \leq 1$. The table also indicates the notation $n_i, i = 1, \ldots, 4$ we use for the number of occurrences of the four pairs $00, 11, 01, 10$ in a sample of size n. It should be clear that n_3 and n_4 play no real role in the problem, since algebraically it must be the case that $|n_3 - n_4| \leq 1$, so for a large sample $n_3 \approx n_4 \approx \frac{1}{2}(n - n_1 - n_2)$.

x_{k-1}	x_k	$p^\theta(x_k \mid x_{k-1})$	$p^\mu(x_k \mid x_{k-1})$	# in sample
0	0	$1-\theta$	$1-\mu$	n_1
1	1	$1-\theta f$	μ	n_2
0	1	θ	μ	n_3
1	0	θf	$1-\mu$	n_4

TABLE A.1. Markov chain transition probabilities, $f = (1-\mu)/\mu$. The sample size is $n = n_1 + n_2 + n_3 + n_4$.

For any k the probability mass distribution of Y_k is $m(x) = 1 - \mu + (2\mu - 1)x$, $x = 0, 1$, while for $n > 0$ the distribution of $(Y_0, \ldots, Y_n)$ is

$$p_n^\theta(x_0, \ldots, x_n) = m(x_0) \prod_{k=1}^{n} p^\theta(x_k \mid x_{k-1}).$$

When $\theta = \mu$ the Y_k are i.i.d. with joint distribution $p_n^\mu(x_0,\dots,x_n) = \prod_0^n m(x_k)$ so, referring to Table A.1, the likelihood ratio $\mathrm{LR}_n = p_n^\theta/p_n^\mu$ is given by

$$\begin{aligned}\mathrm{LR}_n^\theta(x_0,\dots,x_n) &= (1-\theta)^{n_1}(1-\theta f)^{n_2}\theta^{n_3}(\theta f)^{n_4}(1-\mu)^{-(n_1+n_4)}\mu^{-(n_3+n_2)}\\ &= \{(1-\theta)^{n_1}(1-\theta f)^{n_2}\theta^{(n_3+n_4)}\}\{f^{n_4}(1-\mu)^{-(n_1+n_4)}\mu^{-(n_3+n_2)}\}\\ &= \{(1-\theta)^{n_1}(1-\theta f)^{n_2}\theta^{n-n_1-n_2}\}\{f^{n_4}(1-\mu)^{-(n_1+n_4)}\mu^{-(n_3+n_2)}\}\end{aligned}$$

and of course $\mathrm{LR}_n^\mu \equiv 1$. The log likelihood ratio is therefore

$$\mathrm{LLR}^\theta(n_1,\dots,n_4) = L^\theta(n_1,n_2) + M(n_1,\dots,n_4)$$

where

$$L^\theta(n_1,n_2) = n_1\log(1-\theta) + n_2\log(1-\theta f) + (n-n_1-n_2)\log\theta \tag{A.1}$$

and $M = \mathrm{LLR}^\theta - L^\theta$ does not depend on θ.

Proposition A.1. *For $\beta \ge \frac12$, the maximum likelihood estimator is given by*

$$\hat\theta(\bar n_1,\bar n_2) = \frac{1}{2f}\left(\bar n_{-2} + f\bar n_{-1} - \sqrt{(f-c_1)^2 + 4f(c_1-c_2)}\right) \tag{A.2}$$

where $\bar n_{-i} = (n-n_i)/n = 1-\bar n_i$.

Proof. To compute the maximum likelihood estimate we maximize L^θ over θ. We have

$$\begin{aligned}\frac{\partial L^\theta}{\partial\theta} &= -\frac{n_1}{(1-\theta)} - \frac{fn_2}{(1-\theta f)} + \frac{n-n_1-n_2}{\theta}\\ &= \frac{Q(\theta)}{\theta(1-\theta)(1-\theta f)}\end{aligned}$$

where

$$Q(\theta) = fn\theta^2 - (n-n_2+(n-n_1)f)\theta + (n-n_1-n_2). \tag{A.3}$$

The discriminant of Q is

$$\begin{aligned}D &= (n-n_2+(n-n_1)f)^2 - 4fn(n-n_1-n_2)\\ &= n^2[(f+c_1)^2 - 4fc_2],\end{aligned}$$

where

$$c_1 = 1 - \frac{n_2+fn_1}{n}, \qquad c_2 = 1 - \frac{n_1+n_2}{n}.$$

Under our standing assumption $\mu \ge \frac12$ we have $f \le 1$ and hence $c_1 \ge c_2$. Now D can be expressed as

$$D = n^2[(f-c_1)^2 + 4f(c_1-c_2)],$$

showing that $D \ge 0$ whatever the values of n_1, n_2. Taking into account that $Q(0), Q(1/f) > 0$ and $Q(1) < 0$ we easily see that Q has a root in each of the intervals $(0,1),(1,1/f)$ so the maximizing $\hat\theta \in (0,1)$ is the smaller of the two roots, which is given by (A.2). □

Proposition A.2. *In any model $\mathfrak{H}_{\mu,\theta_0}$ with $\mu \in [\frac12,1]$, $\theta_0 \in [0,1]$, the estimator $\hat\theta$ is consistent, i.e. $\hat\theta(\bar n_1,\bar n_2) \to \theta_0$ a.s. as $n\to\infty$.*

Proof. Associated to the chain Y_k is the 4-state Markov chain $Y_k, k = 1, 2, \ldots$ where Y_k takes the values $1, 2, 3, 4$ respectively when $(Y_{k-1}, Y_k) = (0,0), (1,1), (0,1), (1,0)$. The transition matrix for this chain is

$$\mathbf{P} = \begin{bmatrix} 1-\theta_0 & 0 & \theta_0 & 0 \\ 0 & \theta_0' & 0 & 1-\theta_0' \\ 0 & \theta_0' & 0 & 1-\theta_0' \\ 1-\theta_0 & 0 & \theta_0 & 0 \end{bmatrix}$$

Consider first the case $\theta_0 \in (0,1)$. Then the chain Y_k is irreducible and recurrent and consequently has a unique stationary distribution $\mathbf{m}$ characterized by the property that $\mathbf{m}'(\mathbf{I}-\mathbf{P}) = 0$, where $\mathbf{I}$ is the 4×4 identity matrix. This system of equations is readily solved to give

$$\mathbf{m} = \begin{bmatrix} (1-\theta_0)(1-\mu) \\ \mu - (1-\mu)\theta_0 \\ \theta_0(1-\mu) \\ \theta_0(1-\mu) \end{bmatrix},$$

where we have substituted for θ_0' from (5.4). The numbers n_i introduced above are simply the numbers of visits to state i by the chain Y in a sample of length n. Since Y is recurrent,

$$\lim_{n\to\infty} \frac{n_i}{n} = \mathbf{m}_i \quad \text{a.s.}, \quad i = 1, \ldots, 4. \tag{A.4}$$

The quadratic form Q of (A.3) can be written as

$$\frac{1}{n}Q(\theta) = f\theta^2 - (1 - \bar{n}_2 + (1-\bar{n}_1)f)\theta + (1 - \bar{n}_1 - \bar{n}_2). \tag{A.5}$$

If we substitute $\bar{n}_i = \mathbf{m}_i$, $i = 1, 2$ we find $Q(\theta_0) = 0$ and hence that $\hat{\theta}(\mathbf{m}_1, \mathbf{m}_2) = \theta_0$. Now $\hat{\theta}(\bar{n}_1, \bar{n}_2)$ is a continuous function of the two parameters, and hence in view of (A.4) we have $\lim_{n\to\infty} \hat{\theta}(\bar{n}_1, \bar{n}_2) = \theta_0$ a.s.

We now consider the cases $\theta_0 = 0, 1$. When $\theta_0 = 0$, $Y_k = Y_0$ for all k, so either $n_1 = n, n_2 = 0$ or $n_1 = 0, n_2 = n$ giving, from (A.1), values of L^θ equal to $n\log(1-\theta)$ or $n\log(1-\theta f)$ respectively. In either case, L^θ is maximized at $\theta = 0 = \theta_0$.

The case $\theta_0 = 1$ is a little more tricky. Here $Y_{k-1} = 0 \Rightarrow Y_k = 1$, so $n_1 \equiv 0$. The sample path consists of strings of ones separated by single zeros. The probability of flipping from 1 to 0 is f, so the mean length of a string of ones is $1/f$. Each flip from 1 to 0 and back adds 1 to n_3 and to n_4, and each string of ones of length m adds $m-1$ to n_2. So the mean growth rates in $n_3 + n_4$ and in n_2 are in the ratio $2 : (1/f) - 1 = (2\mu - 1)/(1-\mu)$, implying that, loosely speaking, the fraction of the time spent growing n_2 is $((2\mu-1)/(1-\mu))/((2\mu-1)/(1-\mu) + 2) = 2\mu - 1$. We conclude that

$$\lim_{n\to\infty} \bar{n}_2 = 2\mu - 1 \quad \text{a.s.}$$

At the limiting value,

$$\frac{1}{n}L^\theta = (2\mu - 1)\log(1-\theta f) + 2(1-\mu)\log\theta$$

and the derivative with respect to θ is

$$-\frac{\mu(1-f)f}{1-\theta f} + \frac{2(1-\mu)}{\theta}.$$

This is equal to $+\infty$ at $\theta = 0$ and is finite and decreasing for $\theta > 0$. Its value at $\theta = 1$ is $1 - \mu > 0$, and we conclude that the maximum occurs at $\theta = 1$. A simple continuity argument now shows that $\lim_{n\to\infty} \hat{\theta}(0, \bar{n}_2) = 1 = q_0$ a.s. □

References

Abdous, B. and B. Remillard (1995). Relating quantiles and expectiles under weighted-symmetry. *Ann. Inst. Statist. Math.* **47**, 371–384.

Acerbi, C. and B. Szekely (2014, November). Backtesting expected shortfall. *Risk Magazine* **16**.

Artzner, P., F. Delbaen, M. Eber, and D. Heath (1999). Coherent measures of risk. *Mathematical Finance* **9**, 203–228.

Basel Committee on Banking Supervision (2013). Fundamental Review of the Trading Book, a revised market risk framework. 2nd Consultative Paper. Bank for International Settlements, Basel.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedacicity. *Econometrics* **51**, 307–327.

Campbell, J., A. Lo, and A. MacKinley (1990). *The Econometrics of Financial Markets*. Princeton University Press.

Cervera, J. and J. Muñoz (1996). Proper scoring rules for fractiles. In J. Bernardo, J. Berger, P. Dawid, and A. F. M. Smith (Eds.), *Bayesian Statistics*, Volume 5. Oxford University Press.

Christoffersen, P. (1998). Evaluating interval forecasts. *Int. Econ. Rev.* **39**, 841–862.

Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance* **1**, 223–236.

Cont, R., R. Deguest, and G. Scandolo (2010). Robustness and sensitivity analysis of risk measurement procedures. *Quantitative Finance* **10**, 593–606.

Cox, D., D. Hinkley, and O. Barndorff-Nielsen (Eds.) (1996). *Times Series Models in Econometrics, Finance and Other Fields*. Chapman & Hall.

Dawid, A. P. (1984). Present position and potential developments: some personal views. Statistical theory: the prequential approach (with discussion). *J. Roy. Statist. Soc. A* **147**, 278–292.

Dawid, A. P. (1986). Probability forecasting. In S. Kotz, N. Johnson, and C. Read (Eds.), *Encyclopaedia of Statistical Sciences*, Volume **7**, pp. 210–218. Wiley-Interscience.

Dawid, A. P. and V. Vovk (1999). Prequential probability: principles and properties. *Bernoulli* **5**, 125–162.

Dempster, M. A. H. and V. Leemans (2006). An automated FX trading system using adaptive reinforcement learning. *Expert Systems with Applications* **30**, 543–552.

Diebold, F. X. and Mariano, R. (1995). Comparing predictive accuracy. *Journal of Business Economics & Statistics* **13**, 253–263.

Diebold, F., T. Gunther, and A. Tay (1998). Evaluating density forecasts with applications to financial risk management. *Int. Econ. Rev.* **39**, 863–883.

Dudley, R. (1989). *Real Analysis and Probability*. Wadsworth & Brooks/Cole, Pacific Grove CA.

Embrechts, P. and M. Hofert (2014). Statistics and quantitative risk management for banking and insurance. *Annual Review of Statistics and its Applications* **1**, 493–514.

Embrechts, P., C. Klüppelberg, and T. Mikosch (1997). *Modelling Extremal Events: for Insurance and Finance*. Springer.

Engle, R. and S. Manganelli (2004). CAViaR. *Journal of Business and Economic Statistics* **22**, 367–381.

Fissler, T. and J. F. Ziegel (2015, March). Higher-order elicitability and Osband's principle. Available at `arXiv:1503.08123v1`.

Gibbons, J. D. and S. Chakraborti (2010). *Nonparametric Statistical Inference* (5th ed.). Chapman & Hall/CRC.

Gneiting, T. (2011). Making and evaluating point forecasts. *J. Amer. Statist. Assoc.* **106**, 746–762.

Gneiting, T., F. Balabdaoui, and A. Raftery (2007). Probabilistic forecasts, calibration and sharpness. *J. R. Statist. Soc. B* **69**, 243–268.

Gneiting, T. and A. Raftery (2007). Strictly proper scoring rules, prediction and estimation. *J. Amer. Statist. Assoc.* **102**, 359–378.

Gripenberg, G. and I. Norros (1996). On the prediction of fractional Brownian motion. *J. Applied Prob.* **33**, 400–410.

Haldane, A. (2012, August). The dog and the frisbee. Speech, available at `www.bankofengland.co.uk/Home/Publications/Speeches and Articles`.

Hall, P. and C. Heyde (1980). *Martingale Limit Theory and its Application*. Academic Press, NY.

Holzmann, H. and M. Eulert (2014). The role of the information set for forecasting—with applications to risk management. *Annals of Applied Statistics* **8**, 595–621.

Joliffe, I. and D. Stephenson (Eds.) (2003). *Forecast verification: a Practitioner's Guide*. John Wiley.

Kou, S., X. Peng, and C. Heyde (2013). External risk measures and Basel accords. *Mathematics of Operations Research* **38**, 393–417.

Kuester, K., S. Mittnik, and M. Paolella (2006). Value-at-risk prediction: a comparison of alternative strategies. *J. Finan. Econometrics* **4**, 53–89.

Lai, T. L., D. Shen, and S. Gross (2011). Evaluating probability forecasts. *Ann. of Statistics* **39**, 2356–2382.

Lambert, N., D. Pennock, and Y. Shoham (2008, July). Eliciting properties of probability distributions. In *Proc. 9th ACM Conference on Eelectronic Commerce*, Chicago IL.

Mandelbrot, B. (1997). *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk*. Springer.

Mandelbrot, B. and H. Taylor (1967). On the distribution of stock price differences. *Operations Research* **15**, 1057–1062.

Mitchell, J. and K. Wallace (2011). Evaluating density forecasts: forecast combinations,model mixtures, calibration and sharpness. *Journal of Applied Econometrics* **26**, 1023–1040.

Osband, K. and S. Reichelstein (1985). Information-eliciting compensation schemes. *J. Public Econ.* **27**, 107–115.

Popper, K. (2002). *The Logic of Scientific Discovery*. Routledge. [Originally published as *Logik der Forschung*, Vienna 1935.].

Rockafellar, R. T. and S. Uryasev (2000). Optimization of conditional value-at-risk. *Risk* **2**, 21–41.

Rockafellar, R. T. and S. Uryasev (2002). Conditional value-at-risk for general loss distributions. *J. Banking and Finance* **26**, 1443–1471.

Rosenblatt, M. (1952). Remarks on a multivariate transformation. *Ann. Math. Stat.* **23**, 470–472.

Savage, L. (1971). Elicitation of personal probabilities and expectations. *J. Amer. Statist. Assoc.* **66**, 783–801.

Smith, R. C. (2014). *Uncertainty Quantification: Theory, Implementation, and Applications*. SIAM Publications.

Steinwart, I., C. Pasin, R. Williamson, and S. Zhang (2014). Elicitation and identification of properties. In M. F. Balcan and C. Szepesvar (Eds.), *Proceedings of the 27th Conference on Learning Theory 2014*, Volume 35 of *JMLR Workshop and Conference Proceedings*, pp. 482–526.

Warner, T. T. (2010). *Numerical Weather and Climate Prediction*. Cambridge University Press.

Williams, D. (1991). *Probability with Martingales*. Cambridge University Press.

Ziegel, J. F. (2014). Coherence and elicitability. *Mathematical Finance*. (Published online September 2014).

Department of Mathematics, Imperial College London
E-mail address: `mark.davis@imperial.ac.uk`